When Product Markets Become Collective Traps: The Case of Social Media*

Leonardo Bursztyn† **Benjamin Handel**‡ **Rafael Jiménez-Durán**§ **Christopher Roth**¶

November 6, 2023

Abstract

Individuals might experience negative utility from not consuming a popular product. For example, not using social media can lead to social exclusion or not owning luxury brands can be associated with having a low social status. We show that, in the presence of such spillovers to non-users, standard measures that take aggregate consumption as given fail to appropriately capture welfare. We propose a new methodology to measure welfare that accounts for these consumption spillovers, which we apply to estimate the consumer surplus of two popular social media platforms, TikTok and Instagram. In large-scale, incentivized experiments with college students, we show that, while the standard welfare measure suggests a large and positive surplus, our measure accounting for consumption spillovers indicates a negative surplus, with a large share of active users deriving negative utility. We also shed light on the drivers of consumption spillovers to non-users in the case of social media and show that, in this setting, the "fear of missing out" plays an important role. Our framework and estimates highlight the possibility of product market traps, where large shares of consumers are trapped in an inefficient equilibrium and would prefer the product not to exist.

Keywords: Welfare; Consumption Spillovers; Collective Trap; Coordination; Product Market Traps; Social Media.
JEL Classification: D83, D91, P16, J15

*We thank Daron Acemoglu, Eliana La Ferrara, Matthew Gentzkow, Alex Imas, Ro'ee Levy, Matthew Lowe, Pietro Ortoleva, Ricardo Perez-Truglia, Aakaash Rao, Gautam Rao, Jean Tirole, as well as many seminar participants for useful comments. Caroline Boone, Pietro Ducco, Paul Grass, Anastasiya Nebolsina, Ewan Rawcliffe, Christoph Ziegler, and especially Beyza Gulmezoglu provided excellent research assistance. Roth acknowledges funding from the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy EXC 2126/1-390838866. The research described in this article was approved by the University of Chicago Social and Behavioral Sciences Institutional Review Board and pre-registered in AsPredicted (#137878, #142247, and #144630).

†University of Chicago and NBER, bursztyn@uchicago.edu

‡University of California Berkeley and NBER, handel@berkeley.edu

§Bocconi University and Chicago Booth Stigler Center, rafael.jimenez@unibocconi.it

¶University of Cologne, Max Planck Institute for Collective Goods, briq, CESifo, and CEPR, roth@wiso.uni-koeln.de

1 Introduction

Much of consumption is highly social in nature. In many contexts, the individual utility from consuming a product or service increases as more people consume it. Going to a concert or dinner with friends is more enjoyable than going alone. Yet, consumption can also negatively affect others (Frank, 2005). Indeed, the literature on conspicuous consumption and positional externalities (Frank, 1985a; Bursztyn et al., 2018; Imas and Madarász, 2022) has highlighted that one's utility can be negatively impacted by others' incomes or consumption, for instance, as a result of social comparisons (Bottan and Perez-Truglia, 2022; Clark and Oswald, 1996; Cullen and Perez-Truglia, 2022; Luttmer, 2005; Perez-Truglia, 2020).

These social forces play a vital role in the context of social media. For a given platform, a larger number of users may increase the benefits of joining, by expanding the network of individuals available for interaction. Beyond that, the network of users may also affect the utility of potential non-users. Such consumption spillovers to non-users can be driven by mechanisms such as social exclusion or a fear of missing out (Gupta and Sharma, 2021). As the total number of platform users increases, marginal users may participate because they want to avoid these negative non-user spillovers but still have negative overall utility from the platform's existence.

In this paper, we show with a simple conceptual framework and experimental evidence that, in the presence of such consumption spillovers to non-users, the standard consumer surplus measure does not appropriately capture individual welfare. The standard measure—the *individual consumer surplus*—considers one's valuation of consuming the product given the level of others' consumption. As such, it ignores the utility that someone derives from others' consumption of a product when they themselves do not consume it, which we call the *non-consumer surplus*. The assumption that *non-consumer surplus* is zero is plausible for many products but is likely violated for others, e.g., for social media platforms, where non-consumers may derive negative utility from being excluded from interacting with users, for example. In this scenario with a negative non-consumer surplus, the standard individual welfare measure overstates the total welfare associated with the product because it uses an incorrect outside option.[1] Moreover, we show that negative non-

[1] When non-consumer surplus is non-zero, the relevant outside option for calculating welfare is the non-existence of the product market. Also, note that negative non-consumer surplus is an externality to *non-users* only: the product creates the externality to non-users and the only way for individuals to avoid it is precisely by consuming the product. This feature distinguishes it from traditional externalities where

consumer surplus can give rise to *product market traps*: coordination failures where some consumers are trapped in an inefficient equilibrium and would prefer the product not to exist. In such traps, users' utility is negative but would have been even more negative had they not used the platform, which is why they continue using it. This coordination failure can arise from social forces even with fully rational expectations and without behavioral frictions, such as a lack of self-control and naivete.

Guided by our simple framework, we propose a novel methodology to measure consumer welfare in the presence of both network effects and consumption spillovers to non-users and apply it to the welfare analysis of social media platforms. We implement our methodology in pre-registered large-scale online experiments with more than 1,000 students from various colleges in the US. We focus on two prominent social media platforms, TikTok and Instagram, that have been the subject of concern, among other reasons, due to their potential adverse effects on mental health (Faelens et al., 2021).[2]

In the experiment, we employ standard tools to measure consumer welfare, leveraging an incentivized Becker-DeGroot-Marschak (BDM) mechanism (Becker et al., 1964), which we implement using an iterative multiple price list following best practices (Berkouwer and Dean, 2022). The experiment proceeds in three main steps: In step 1 (*Valuation Keeping Network*), we measure individual-level willingness to accept (WTA) to deactivate one's social media for four weeks while keeping constant others' social media consumption. This step provides us with the standard measure of individual consumer surplus. In steps 2 and 3, we plausibly reduce the size of our respondents' networks by presenting the possibility of a large-scale deactivation where all participating students at their university deactivate their accounts. Participants are truthfully told that this large-scale deactivation will be conducted if we recruit two-thirds of students at their university. To measure network effects, we measure individual WTA conditional on all participating students being asked to deactivate their account in exchange for monetary compensation in step 2 (*Valuation Removing Network*). Finally, in step 3 (*Product Market Valuation*), we measure welfare accounting for both consumer and non-consumer surplus. To do so, we elicit individuals' preferences over the deactivation of the social media accounts of all participating students, including themselves. In particular, we measure whether individuals are willing to forego

both users and non-users are affected.

[2] We measured TikTok's welfare in a survey conducted in July 2023, and Instagram's welfare in a survey conducted in August and September 2023. Both surveys are virtually identical, except that in the second survey we added more questions and clarified some of the instructions. We describe the differences in Section 3.2 and present the full set of instructions in Appendix C.

payment or instead require a payment to deactivate all participating students' accounts.

Our main results highlight the importance of accounting for consumption spillovers to non-users. Our individual-level elicitation estimates indicate an average individual consumer surplus of \$59 and \$47 for TikTok and Instagram, respectively, with 93% and 86% of users deriving positive welfare from the product. These findings are in the ballpark of estimates in the literature (Mosquera et al., 2020; Allcott et al., 2020) and indicate that users require substantial compensation to stop using social media when others in their network keep using it.

We next turn to product market valuation, our preferred measure of welfare that accounts for consumption spillovers to non-users. Product market surplus is significantly lower compared to individual consumer surplus for both TikTok ($p < 0.01$) and Instagram ($p < 0.01$). Users have an average willingness to pay (WTP), rather than a willingness to accept, of \$28 and \$10 to have others, including themselves, deactivate TikTok and Instagram, respectively. Sixty-four percent and forty-eight percent of active TikTok and Instagram users experience negative welfare from the product's existence. Participants who do not have an account have an average WTP of \$67 and \$39 to have others deactivate TikTok and Instagram, respectively. Overall, our evidence shows the existence of a *social media trap* for a large share of consumers, who find it individually optimal to use the product even if they derive negative welfare from it.

Finally, we present our estimates of network effects by comparing the valuation removing the network against the valuation keeping the network. This measure indicates an average willingness to accept of \$40 and \$36, constituting a significant drop of approximately 33% and 24% compared to the *Valuation Keeping Network*, for TikTok and Instagram, respectively. Moreover, the fraction with positive welfare drops to approximately 70% and 69% of users for TikTok and Instagram, respectively. As we argue with our framework, this drop provides evidence that network effects are positive and quantitatively significant, consistent with canonical theoretical frameworks (Rohlfs, 1974; Katz and Shapiro, 1985).

To ensure high levels of understanding, we restrict our main analysis to respondents who pass several attention checks and do not regret their choices, though our results are robust to including regretters and inattentive respondents. Moreover, we confirm our findings using a hypothetical qualitative question in which we ask respondents whether they would prefer to live in a world without the social media platform. Indeed, most respondents and a large share of users in our samples would prefer to live in a world without TikTok and Instagram, respectively. Similarly, another hypothetical question reveals that respondents favored the option of everyone deactivating their accounts over only deactivating their own account or no one deactivating their accounts, for both Instagram and TikTok.

One possible concern with our empirical design is that respondents may not think that it is likely that we will actually conduct the large-scale deactivation study. However, the perceived likelihood that the large-scale Instagram deactivation study will be implemented is high at approximately 45%. Moreover, for respondents deeming the large-scale deactivation study more likely, the estimated product market surplus is even more negative, suggesting that our elicitation provides a conservative estimate of how negative the product market surplus is. More broadly, given that even in the case of the large-scale deactivation study not all users would deactivate their account, our study plausibly identifies lower bounds for the size of negative product market surplus.

One conjecture is that the drop in welfare from the individual to the product market valuation could be fully driven by factors such as a "repugnance" (Roth, 2007) towards digital products, animus against big tech companies, or a distaste for others spending time on their phone. To rule out this possibility, we conduct an experiment with an identical design but with a product that has plausibly less pronounced negative consumption spillovers to non-users: navigation and maps smartphone apps.[3] For these apps, our estimates of product market surplus remain positive, large, and highly significant. Besides elucidating the relative importance of different mechanisms, the positive product market valuation for maps also suggests that the negative product market valuation we document for TikTok and Instagram is not driven by mechanical factors such as the way we frame our elicitation.

The wedge between individual consumer surplus and product market surplus highlights an important role of consumption spillovers to non-users. To shed light on the motives behind active users' preferences for living in a world without their social media platform, we ask them an open-ended question on why they still use the platform. This data indicates that the fear of missing out is the most prevalent motive for both TikTok and Instagram. To provide even more direct evidence on the nature of consumption spillovers to non-users, we ask respondents another open-ended question about their feelings if they were the only ones who had to deactivate their accounts and everyone else kept using them. A large fraction of active Instagram and TikTok users express negative feelings, particularly the fear of missing out. Paired with our main estimates, the evidence of these underlying mechanisms supports the notion that accounting for non-consumer surplus is crucial to assessing the welfare effects of social media platforms. These consumption spillovers to

[3]Hereafter referred to simply as "Maps."

non-users may arise from anticipated social exclusion that would actually occur in case of deactivation or could arise from misperceptions (Bursztyn and Yang, 2022) or other psychological biases.

Our estimates thus far highlight the possibility of product market traps on social media, where consumers are trapped in an inefficient equilibrium and have a preference for the product not to exist. To provide suggestive evidence on the existence of product market traps in other domains, we fielded online surveys with US consumers in the context of luxury goods, where positional externalities are a plausible driver of consumption spillovers to potential non-consumers. We use our validated hypothetical questions to document that among respondents who owned luxury brands that they themselves bought (e.g., Gucci, Versace, Rolex), 44% preferred to live in a world without any of those brands altogether. Among respondents not owning such brands, the fraction preferring to live in a world without them is 69%. Taken together, this evidence underscores that consumption spillovers to non-users occur across many different markets. While the literature on luxury goods and conspicuous consumption has emphasized the negative externalities these goods impose on non-consumers (Frank, 1985a, 2000, 2012), our evidence highlights that even large shares of consumers of these products would prefer them not to exist.

Moreover, product market traps lead to a situation where the existence of a product is harmful to consumers, which can manifest not only as excessive consumption by users but also as the production of an excessive number of product variations or vintages (Pesendorfer, 1995). To measure whether consumers consider the frequency of vintages of some products excessive, we asked respondents whether they prefer to live in a world where Apple releases the iPhone every year or every other year. Among iPhone owners, a striking 91% of respondents indicated that they would prefer Apple to release the iPhone every other year rather than every year. Among respondents not owning the iPhone, this fraction was even larger, at 94%. This evidence suggests that consumers consider the number of product variations or vintages of the iPhone as excessive and thus harmful to consumer welfare.

One implication of our framework is that producers have incentives to use technologies or marketing campaigns that decrease non-consumer surplus – increasing the cost of not consuming the product. Consistent with this, large tech companies commonly use tools that might decrease non-consumer surplus, such as increasing the salience of being a non-consumer or tying together messaging apps and social media platforms. An example of such technology is the case of the “green bubble” messages on iPhones, which make it salient for iPhone users when they exchange text messages with non-iPhone users. The social stigma arising from this green bubble culture has received widespread attention in the mass media.[4] In this case, green bubbles might increase demand for the iPhone, not because it improves the intrinsic value of the product, but rather because it avoids the social stigma cost of not using an iPhone.

More broadly, our findings challenge the standard revealed-preference argument that the mere existence of a product implies positive welfare for its consumers, even if they have rational expectations and in the absence of cognitive distortions. Indeed, we provide evidence of a product that is consumed by a large share of individuals, even when it creates negative welfare for many of them. This finding suggests a heightened need for regulators to assess whether different products create traps for consumers and, potentially, diminish competition between platforms (Tirole, 2023).

Our analysis relates to a large literature on consumption spillovers and, in particular, "bandwagon" effects, positional externalities, and status goods (Leibenstein, 1950; Pesendorfer, 1995; Frank, 1985a,b, 2000, 2012; Becker, 1991; Bagwell and Bernheim, 1996; Heffetz and Frank, 2011). Previous empirical work has demonstrated the importance of peer effects in consumption (Kuhn et al., 2011) and financial decisions (Bursztyn et al., 2014), documented a large demand for status goods (Bursztyn et al., 2018), and a higher willingness to pay for more exclusive products (Imas and Madarász, 2022). We contribute to this literature by providing the first empirical evidence on estimates of consumer welfare accounting for non-consumer surplus.[5] Our methodology for estimating welfare also contributes to the behavioral public economics literature, which examines public policy and welfare in the presence of non-standard preferences (Bernheim, 2016; Bernheim and Taubinsky, 2018; Butera et al., 2022; Ambuehl et al., 2021).[6]

Our paper also speaks to work assessing the welfare created by social media (Brynjolfsson et al., 2019; Mosquera et al., 2020; Allcott et al., 2020, 2022; Brynjolfsson and Oh, 2023; Brynjolfsson et al., 2023). The papers in this space measure consumer surplus by either taking the aggregate level of consumption as given or assuming that consumption spillovers to non-users are zero. Existing work finds large user valuations for social

[4]For popular press coverage on the "green bubble culture": See "Why Apple's iMessage Is Winning: Teens Dread the Green Text Bubble" Higgins, Tim. *The Wall Street Journal*, January 8, 2022.

[5]DellaVigna et al. (2012) find that door-to-door charity fundraising campaigns on average lower the utility of potential donors due to social pressure costs.

[6]Our paper also relates to a literature on contingent valuation in environmental and resource economics, which shows that preferences can be reverted depending on whether goods are valuated individually or jointly (Alevy et al., 2011).

media platforms consistent with the large amount of time spent on these platforms (2.5 hours daily for Americans (Kemp, 2022)), while at the same time documenting that the expansion and use of these platforms can harm individual well-being, particularly mental health (Braghieri et al., 2022). Our results on the switch in signs of consumer welfare after accounting for non-consumer surplus help reconcile these seemingly contradictory findings and paint an integrated and more pessimistic picture of the welfare effects of social media. Additionally, we provide the first incentivized evidence of network effects on welfare in the context of social media, which has proven difficult aside from existing hypothetical estimates (Benzell and Collis, 2022).

We also contribute to a long-standing literature in industrial organization that models consumer choice in the presence of network effects (Rohlfs, 1974; Katz and Shapiro, 1985; Farrell and Klemperer, 2007; Rochet and Tirole, 2003). Jullien et al. (2021) provide an overview of papers that estimate network effects and note (i) the myriad challenges in doing so with administrative data and (ii) the close linkage between the identification challenges in this literature and the broader literature on peer effects in demand (Manski, 1993). Our work differs from this literature in a few key ways. First, a standard procedure in this literature is to normalize the utility from not using a product to zero, effectively ruling out consumption spillovers to non-users.[7] We relax this assumption and allow for both traditional user network effects and consumption spillovers to non-users to flexibly change (and co-vary) as the product user base changes. Second, we develop an experimental framework to elicit the magnitude of network effects and consumption spillovers to non-users, while the literature typically uses administrative data together with instruments to separately identify network effects. We simultaneously identify and quantify both positive network effects for product users and negative consumption spillovers to non-users. Finally, the literature has pointed to coordination failures that arise in the presence of network effects in cases where one firm becomes dominant despite not being the most efficient supplier (Farrell and Saloner, 1985; Farrell and Klemperer, 2007).[8] While this coordination failure occurs in the presence of multiple competing platforms and consumption spillovers among product users, our work highlights the possibility of *product market traps*: a coordina-

[7] One exception is Bhattacharya et al. (2023) who also relax this assumption and show that welfare effects are not point identified in models with consumption spillovers. They apply their model to the evaluation of the welfare effects of bed nets in a discrete-choice econometric framework with a focus on health externalities arising from contagious diseases. Another exception is the literature on informational externalities (Choi et al., 2019; Acemoglu et al., 2022), where non-users are affected by data disclosure of users.

[8] Coordination failures may also explain the persistence of harmful social norms (Ferrara et al., 2023).

tion failure that can arise even with a single platform due to the presence of consumption spillovers to non-users that lock consumers into using the product.[9]

This paper proceeds as follows: Section 2 provides the conceptual framework for analyzing welfare effects in the presence of consumption spillovers to potential non-users and introduces the notion of *product market traps*. Section 3 provides the empirical design and results for individual and product market surplus. Section 4 presents survey evidence on product market traps in the context of luxury and vintage goods. Finally, Section 5 discusses the policy implications of our findings.

2 Conceptual Framework: Product Market Traps

Setup. A continuum of individuals simultaneously decide whether to buy an indivisible product with a fixed price p. Individual i derives quasilinear utility $u_i(x_i, X_i^e)$ where $x_i \in \{0, 1\}$ is their own consumption of the product and X_i^e is their belief about the fraction of other individuals expected to consume it.

We leverage the presence of X_i^e in an individual's utility to model two distinct phenomena that relate group consumption to that utility. First, we allow for both consumption spillovers—the extent to which utility changes directly in response to others' consumption–and network effects—the degree of complementarity between individual consumption and others' consumption. Social media platforms are a classic example of product network effects, where the presence of many consumers on the platform makes using the platform more attractive for an individual user.[10]

Second, we relax an assumption commonly held in the literature and allow for consumption spillovers among non-users. This means that individual utility from *not consuming* the product can vary with X_i^e. This could be due to, e.g., concerns about product-affiliated status (Frank, 1985a,b; Heffetz and Frank, 2011), due to a fear of missing out (Gupta and Sharma, 2021), or due to exhibiting repugnance towards others' consumption (Roth, 2007).

[9] While we focus on the latter coordination failure, in the presence of multiple platforms both kinds of coordination failures could be present simultaneously. However, the empirical patterns of social media use, paired with our finding of negative welfare among single- and multi-homers (those who use one platform or multiple platforms, respectively), suggest product market traps above and beyond the coordination failure in Farrell and Saloner (1985). Additionally, product market traps are a form of Prisoner's Dilemma, which relates to recent literature documenting forms of prisoner's dilemma in the industry generated by different mechanisms (Cherye and Acquisti, 2022; Sullivan, 2022).

[10] Note that the framework is general enough to allow individuals to get spillovers from both the fraction of users X_i^e or the fraction of non-users $1 - X_i^e$.

Below, we show that the presence of consumption spillovers to non-users generates a wedge between traditional *individual* welfare measures based on revealed preferences conditional on X_i^e and welfare measures that allow for spillovers to non-users.

As a result, we show that standard measures can fail to appropriately capture welfare, and we outline conditions under which a coordination failure arises and traps consumers in an inefficient equilibrium, a product market trap. The mechanism leading to this inefficient equilibrium is not due to the standard network effects problem with multiple platforms (Farrell and Saloner, 1985), but in addition to that problem, arising from consumption spillovers to non-users.

Consumption Spillovers. We define consumption spillovers as a change in individual utility in response to a change in the fraction of other consumers using the product. Specifically, i exhibits positive consumption spillovers from the product when, for $X' > X$ and fixed x_i:

$$u_i(x_i, X') > u_i(x_i, X),$$

with negative consumption spillovers defined analogously.

As is standard in the literature, we normalize to zero the utility that i receives when they believe that no one else consumes the product, $u_i(0,0) = 0$. However, most prior work on network effects, and prior empirical work on social media, does not allow for consumption spillovers to non-users. Thus, prior work typically assumes that for all X and X':

$$u_i(0, X) = u_i(0, X'),$$

which implies that $u_i(0, X) = 0$ for all X. We relax this assumption and allow for consumption spillovers for non-users, i.e., we allow for:

$$u_i(0, X) \neq u_i(0, X').$$

If, e.g., $u_i(0, X') < u_i(0, X)$ when $X' > X$, we say that the product exhibits negative consumption spillovers for non-users. Note that this is an externality to non-users only: the product creates the externality to non-users and the only way to avoid it is precisely by consuming the product.

Given this relaxation, we need to distinguish between network effects and consumption spillovers. We use a definition of direct network effects based on strategic complementarities in consumption. Concretely, individual i's utility exhibits positive network effects when their marginal utility of consumption increases with the consumption of others:

$$u_i(1, X') - u_i(0, X') > u_i(1, X) - u_i(0, X),$$

for $X' > X$. In words, the utility gain from consuming the product is increasing in others' consumption. In the absence of consumption spillovers for non-users, this definition based on strategic complementarities is equivalent to the standard definition of network effects in the literature of having positive consumption spillovers conditional on using the product (e.g., $u_i(1, X') > u_i(1, X)$).

Welfare. A standard measure of individual consumer surplus in the literature compares the utility that i gets relative to their utility when they do not consume the product, given the number of X_i^e others consuming it. We refer to this measure as the *individual* consumer surplus ICS, in the sense that it only accounts for i's individual choice:

$$ICS_i(p, X_i^e) \equiv \begin{cases} u_i(1, X_i^e) - p - u_i(0, X_i^e) & \text{if } \; i \text{ consumes, } u_i(1, X_i^e) - p \geq u_i(0, X_i^e) \\ u_i(0, X_i^e) - u_i(0, X_i^e) = 0 & i \text{ does not consume, } u_i(1, X_i^e) - p < u_i(0, X_i^e). \end{cases}$$

In practice, researchers estimate this measure by eliciting individuals' willingness to accept to give up a product in exchange for a monetary payment, or their willingness to pay to get it, holding constant the others' consumption.

We introduce a welfare measure that allows for consumption spillovers to non-users, which we call the *product market surplus*. This measure compares the utility that i gets relative to their utility when no one consumes the product:

$$PMS_i(p, X_i^e) \equiv \begin{cases} u_i(1, X_i^e) - p - u_i(0, 0) = u_i(1, X_i^e) - p & \text{if } \; i \text{ consumes} \\ u_i(0, X_i^e) - u_i(0, 0) = u_i(0, X_i^e) & \text{if } \; i \text{ does not consume.} \end{cases}$$

The key difference between these measures is the different outside option that each uses. Product market surplus is better suited to measure welfare W from the product's existence and use because it correctly compares the utility that consumers and non-consumers get from a product to their utility in the absence of the product:

$$
\begin{aligned}
W &\equiv \int_i PMS_i(p, X_i^e)\ \mathrm{d}i \\
&= \underbrace{\int_{\text{Consumers}} [u_i(1, X_i^e) - p]\ \mathrm{d}i + \int_{\text{Non-consumers}} u_i(0, X_i^e)\ \mathrm{d}i}_{\text{With the product}} \underbrace{- \int_i u_i(0,0)\ \mathrm{d}i.}_{\text{Without the product}} \qquad (1)
\end{aligned}
$$

In the absence of consumption spillovers, both measures are identical; $ICS_i = PMS_i$, and the product market surplus is simply the sum of individual consumer surpluses. More generally, however, the individual consumer surplus will be biased upwards or downwards depending on whether $u_i(0, X_i^e)$—which we call *non-consumer surplus*—is on average negative or positive, respectively:

$$
\int_i ICS_i(p, X_i^e)\ \mathrm{d}i = \int_i PMS_i(p, X_i^e)\ \mathrm{d}i - \int_i u_i(0, X_i^e)\ \mathrm{d}i.
$$

In other words, the presence of negative or positive consumption spillovers when i does not consume the product determines whether the individual consumer surplus is higher or lower than the product market surplus. For example, when i has a fear of missing out, their individual surplus will be biased upward, as it reflects not only their valuation of the product but also their distaste for being left out when they do not consume it.

Equilibrium. Consider a rational expectations equilibrium which requires consistency between individual choices and expectations. As such, the fraction of consumers X has to be equal to the demand $D(p, X)$ for the product when individual beliefs equal X:

$$
X = D(p, X) \equiv \int_i \mathbb{1}_{\{u_i(1,X) - p \geq u_i(0,X)\}}\ \mathrm{d}i. \qquad (2)
$$

Under standard regularity conditions, the right-hand side of the previous expression is decreasing in p and we can obtain the inverse demand curve $p(X)$. The equilibrium fraction of consumers $X^*(p)$ satisfies $p(X^*) = p$, and need not be unique. As Appendix A details, we can use the introspective equilibrium notion of Akerlof et al. (2023) to select among possible multiple equilibria.

Figure 1 illustrates some key aspects of our framework. The example in the figure assumes uniform and negative consumption spillovers to non-users as well as an equilibrium price of 0, all for illustrative purposes. $D(p, X^*)$ represents the consumer demand curve

Figure 1: Static Welfare Effects Given Market Share X^*

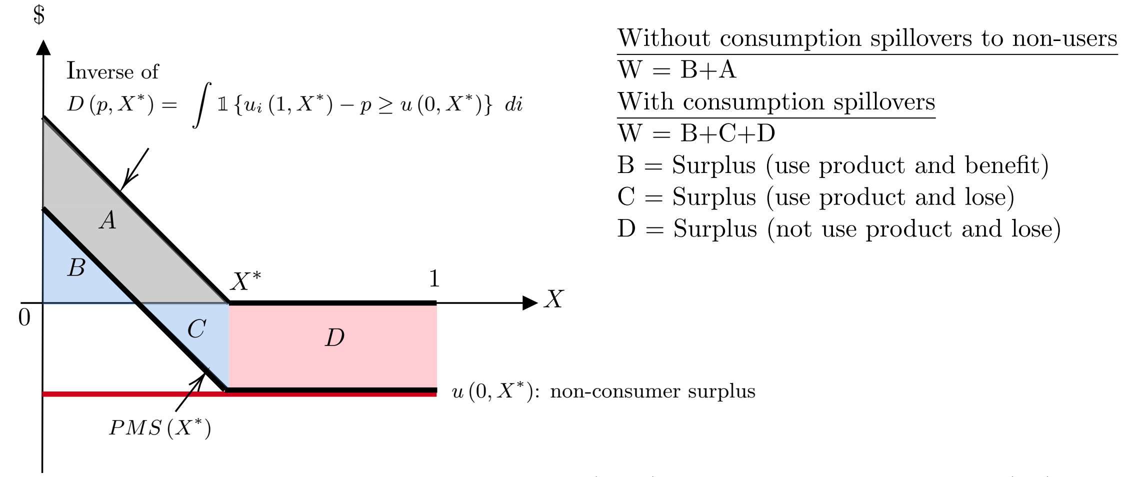

Notes: Figure 1 presents the inverse of the demand curve $D(p, X^*)$, the product market surplus $PMS(X^*)$, and the non-consumer surplus $u(0, X^*)$ conditional on an equilibrium market share X^*. The standard measure of consumer surplus, ignoring consumption spillovers to non-users, is given by areas A+B. Our measure accounting for such spillovers is given by B+C+D.

for the product when the equilibrium usage of the product is X^*. This demand curve thus reflects willingness to pay given any consumption spillovers (positive or negative) with X^* other users. Importantly, in our framework, the demand curve conditional on X^* also considers that a consumer gets negative utility from not consuming the product when others are consuming it. This negative utility is the horizontal red line in the chart.

A classic welfare analysis would estimate the triangle under the demand curve but above the price to determine aggregate consumer surplus. In the simple example above, this classic approach would say that the consumer surplus generated by the product is A + B. Everyone who uses the product benefits from it and everyone who does not use the project gets 0 utility from it. In our framework, the welfare calculation is more subtle. Now, there are non-zero welfare impacts for three distinct groups: (i) users who benefit from the product (ii) users who get negative utility from the product, and (iii) non-users who get negative utility from the product. The welfare impact of an equilibrium with X^* users and a price of 0 is now (i) positive and equal to the area of B for users who benefit from the product (ii) negative and equal to the area of C for users who lose out from the product and (iii) negative and equal to D for non-users who lose out from the product's existence. Thus, in our framework accounting for consumption spillovers to non-users, the estimated welfare is W = B + C + D, instead of A + B in the classical approach.[11] Note that the figure is conditional on an equilibrium use level X^*. In Appendix A, we include details of how our model and empirical framework inform comparative statics related to changes in equilibrium use.

Product Market Traps. Our framework allows for the possibility of a *Product Market Trap*, which is an equilibrium where (i) consumers use the product and (ii) they would be better off if no one consumed it. In a product market trap, individuals would like to coordinate to not consume the product, but they cannot commit to doing so; even if their welfare from using the product is negative, they are "trapped" into using it because others do so. Hence, the standard revealed-preference argument that the existence of a product implies that users benefit from it fails to apply. Specifically, for a given equilibrium, we say that individual i experiences a product market trap when:

(i) i chooses to consume the product: $ICS^i(p, X) > 0$.

(ii) i's welfare is negative: $PMS^i(p, X) < 0$.

Note that these conditions imply that consumer i experiences negative non-consumer surplus ($u_i(0, X) < 0$). In other words, negative spillovers to non-users are a necessary condition to generate product market traps. Appendix A explains how positive network effects and early adopters (users who derive positive utility from joining the platform when no one uses it) can explain how product market traps arise in equilibrium.

Table 1 below illustrates the product market trap and the bias of the individual consumer surplus with two identical consumers when $p = 0$. Consider an initial equilibrium where both use the product: $u(1, 1) > u(0, 1)$. Individually, both consumers are better off consuming the product given that the other person is consuming it, so their individual consumer surplus is non-negative ($ICS_i(0, 1) > 0$). A coordination problem arises if they experience negative utility from consumption ($u(1, 1) < 0$), which means that their product market surplus is negative ($PMS_i(0, 1) < 0$). Note that the coordination problem described here does not occur when negative consumption spillovers for non-users are assumed away, as is commonly done in the literature. As mentioned earlier, it is important

[11] It is important to keep in mind that this graph is an illustrative example and our empirical work will determine the precise way in which accounting for consumption spillovers to non-users changes the estimate of the welfare impact of social media. That empirical work accounts for the precise shape of demand, the heterogeneous distribution of non-user spillovers, and the correlation between the two.

to note that these negative consumption spillovers for non-users are in addition to network effects in product use, which in this case arise if $u(1,1) - u(0,1) > u(1,0) - u(0,0)$.

Table 1: Example of a Coordination Problem with Two Consumers

		Consumer 2	
		Consumes	Does not consume
Consumer 1	Consumes	$u(1,1), u(1,1)$	$u(1,0), u(0,1)$
	Does not consume	$u(0,1), u(1,0)$	0,0

Additionally, producers have incentives to use technologies or marketing campaigns that decrease non-consumer surplus $u_i(0, X)$ – increasing the cost of not consuming the product.[12] An example of such technology is the case of the "green bubble" messages on iPhones, which make it salient for iPhone users when they exchange text messages with non-iPhone users. From Equation (2), it is clear that such a technology would increase demand $D(p, X)$ by reducing $u_i(0, X)$, pushing up the producer's revenues.

Consumption Spillovers and Network Effects. Negative consumption spillovers to non-users mean that a larger network size decreases non-consumer surplus. In other words, the larger the network, the more costly it is to stay out or to leave the network. For this reason, negative consumption spillovers to non-users intensify the strength of network effects because they increase the cost of non-consumption. To see why, we write network effects as:

$$\frac{\partial u_i(1, X)}{\partial X} - \frac{\partial u_i(0, X)}{\partial X}.$$

Note that this expression increases as the term $\partial u_i(0, X)/\partial X$ becomes more negative. In words, more negative consumption spillovers to non-users imply that network effects become more positive, all else equal. Hence, the presence of more negative consumption spillovers to non-users implies that fewer early adopters are needed to transition to a given equilibrium (see Appendix A) – or, alternatively, the same number of early adopters results in larger subsequent product penetration.

The relationship between the network size and non-consumer surplus has important implications for optimal policy. While in models without consumption spillovers a larger

[12] Alternatively, producers can deepen the negative consumption spillovers; i.e., make $\partial u_i(0, X)/\partial X$ more negative. If we assume that the technology has no impact when no one uses the product, then the technology also decreases non-consumer surplus for any X.

network increases individuals' welfare, the relationship between the size of the network and welfare is ambiguous when negative consumption spillovers to non-users are present. As can be seen in Equation (1), a larger network size (holding price constant) has two effects on welfare: it changes product market surplus among consumers and non-consumer surplus among non-consumers. With large enough negative consumption spillovers to non-users, if product market surplus among consumers does not increase enough, a larger network size decreases welfare. In this case, anti-trust law or regulatory action to reduce the size of networks (e.g., a Pigouvian tax that increases with network size, as suggested by Romer (2019)) may be welfare-improving.

3 Measuring Individual and Product Market Surplus

3.1 Sample

College Student Sample. We recruited college students to participate in our experiments through a partnership with College Pulse, a company specialized in recruiting college students for online experiments with a panel of 650,000 college students. We focus on college students for various reasons: First, they are of high policy relevance as they are among the most active on social media.[13] Second, social media usage has been linked to the increasing prevalence of depression among college students (Braghieri et al., 2022).[14] Third, even if other fellow college students might not represent the entire network of friends of our participants (i.e., corresponding to X in our theoretical framework), they constitute a significant subset of students' social networks.[15]

Pre-registration. The pre-registrations include the experimental design, hypotheses, analysis, sample sizes, and exclusion criteria. The pre-registrations of the two data collections can be found on AsPredicted #137878 and #142247.[16]

[13] In an April 2021 study by the Pew Research Center, 84% of adults between the ages of 18 and 29 reported using social media (Pew Research Center, 2021). Among our respondents, 93% stated to had used Instagram at least once in the past month, with 74% using it daily. Meanwhile, 66% mentioned using TikTok at least once in the past month, and 51% did so daily.

[14] In May 2023 the US Surgeon General issued an advisory urging a push to better understand the possible social media "harm to the mental health and well-being of children and adolescents."

[15] Our respondents estimate an average 57% of their mutual friends on Instagram are fellow college students, indicating that the college social network constitutes a majority of our respondents' social networks.

[16] We also pre-registered running an in-person experiment at the University of Chicago at the end of May 2023. However, we failed to recruit the minimum pre-registered number of participants and only managed

TikTok. In July 2023, we recruited 1,936 respondents who began our experiment, out of which 66% had used TikTok in the past month, our measure of activity on the platform. All active users are then asked whether they are willing to participate in the deactivation study. 57% of TikTok users in our sample were willing to provide their handle to participate in the study. Much of this selection does not simply arise from an unwillingness to deactivate their accounts, with 40% of participants mentioning privacy concerns, and 32% mentioning the fear of missing out as motives for not being willing to participate in the study (see Appendix Figure A3). We restrict our sample to respondents aged between 18 and 30 and we exclude those who took the experiment more than once, as identified by their user IDs, as well as respondents who failed any attention checks or regretted their valuations for a second time. Our final sample consists of 595 college students, 291 TikTok users, 304 non-users.

Instagram and Maps To provide evidence for a second social media platform and for another smartphone application that is not a social media platform, we recruited college students who had not taken our TikTok experiment to participate in a second wave in August and September 2023. Respondents were randomly assigned to complete a version of the experiment about (i) Instagram or (ii) Maps (navigation and maps smartphone apps such as Google Maps, Apple Maps, and Waze). All active users are asked for their willingness to participate in the deactivation study, while non-users proceed directly to the practice questions and the *Product Market Valuation*.

We randomize a total of 1,495 and 1,487 respondents into the Instagram and Maps experiments, respectively. Out of those, 93% reported actively using Instagram and 99.7% reported using Maps.[17] All active users are then asked whether they are willing to participate in the deactivation study. 46% of Instagram users in our sample were willing to provide their handle to participate in the study, while 58% of Maps users were willing to participate. As with TikTok, much of this selection is not simply a result of their unwillingness to deactivate their accounts, with 32% and 26% of participants mentioning privacy concerns[18], and 42% and 16% mentioning "not wanting to be without their account while

to collect 11 pre-registered responses from active users who passed the sample inclusion criteria. The very small sample thus makes it difficult to draw meaningful conclusions.

[17] To assess how representative our sample is in terms of social media usage, we compare it to data obtained from the American Trends Panel of the Pew Research Center (2021). To increase comparability with our sample, we filter the data by age and education, to approximate a sample of college students. Specifically, we narrow the data to respondents aged 18-29 and those in the education category of "Some college, no degree." Among respondents in this filtered sample, 54% and 75% reported to use TikTok and Instagram, respectively.

[18] Participation in the deactivation study required respondents to providing their TikTok/Instagram

their friends are still on the platform" as motives for not being willing to participate in the study, for Instagram and Maps, respectively.[19]

Our main sample consists of 230 Instagram users, 25 respondents not active on Instagram, and 252 Maps users. As with TikTok, we restrict our sample to respondents aged between 18 and 30 and exclude – as pre-specified – respondents taking the experiment more than once, failing any attention checks, or regretting any of their final valuations.[20]

3.2 Design

The purpose of the experiment is to measure welfare while accounting for consumption spillovers to non-users. Below we describe the core experimental instructions. The full set of instructions can be found in Appendix C.

TikTok and Instagram. Our main evidence focuses on consumers' valuation of two popular social media platforms, TikTok and Instagram, that have been the subject of concern regarding their impact on individual well-being. While both TikTok and Instagram are social media platforms that focus on visual content, they differ in several key ways. TikTok specializes in short-form video content, often featuring music, dance, and challenges, and utilizes a unique algorithm that prioritizes content discovery, allowing even unknown creators to go viral. Instagram, on the other hand, started as a photo-sharing platform and its discovery mechanisms are more reliant on existing social networks and hashtags, making it generally harder for new creators to gain visibility.

Overview. We now turn to the structure of our experiments, which is also summarized in Figure 2 for the case of the Instagram and Maps experiment.[21] For active social media users, the experiment proceeds in four steps. In step 0, we measure individual-level WTA

handles and submitting screenshots of their phone's usage statistics. For Maps, only the usage statistics were required.

[19] Even among respondents unwilling to participate in our Instagram deactivation study a large share (39%) prefer living in a world without Instagram. 3% respondents unwilling to participate in the maps deactivation study prefer living in a world without maps.

[20] All of our attentive respondents use Maps. As opposed to the case of Instagram and Maps, our TikTok pre-registration did not specify dropping inattentive users or those who regret their choices, but we add these filters to increase comparability across samples. However, results are similar to our preferred specification without these filters (see Appendix Figure A10).

[21] The TikTok experiment has a similar structure but with only one platform. This structure applies to both social media (TikTok and Instagram) and Maps users, but for simplicity, we refer to all these platforms as "social media." We reintroduce the distinction when we talk separately about Maps in section 3.5.1.

to deactivate an example product, an app for a ride-sharing platform. This elicitation considers the individual-level decision conditional on aggregate consumption and is meant to accustom respondents to the instructions. In step 1 (*Valuation Keeping Network*), we measure individual level WTA to deactivate one's social media account for a period of four weeks taking others' social media consumption as given. In steps 2 and 3, we present respondents with the possibility of a large-scale deactivation study where all participating students at their university deactivate their accounts. In step 2 (*Valuation Removing Network*), we measure individual WTA conditional on all participating students being asked to deactivate their account in exchange for monetary compensation. In step 3 (*Product Market Valuation*), we measure individuals' preferences over the deactivation of social media accounts of all participating students, including themselves. In particular, we elicit students' WTP or their WTA to deactivate everyone's account.

Respondents who are not active social media users take a modified version of the experiment. After completing the practice, they proceed to a customized *Product Market Valuation*, where we measure their preferences over the deactivation of social media accounts of all participating students who are active social media users.

Introduction. We inform all respondents that we will conduct a deactivation study in which we will ask students at their university to deactivate their social media accounts for four weeks in exchange for monetary compensation. To enhance the credibility of our deactivation study we inform them that "deactivation studies like this have been conducted in the past (e.g., by Allcott et al. (2020) and Mosquera et al. (2020))." We explain that they can go back to using their account whenever they want, with their content and network unchanged, but they would then forgo any monetary payment. We also tell respondents that, to verify that they deactivate their accounts, we will visit their profiles and require them to upload screenshots of their app usage.[22] To ensure high levels of attention, we inform respondents they will receive an additional bonus payment if they correctly respond to all comprehension questions included in the experiment.

Willingness to Accept Elicitation. The core object of interest in our experiment concerns people's willingness to accept the deactivation of their social media accounts for four weeks. We combine an incentivized BDM elicitation (Becker et al., 1964) with

[22]The use of screenshots of participants' app usage prevents them from substituting between different accounts of the same platform.

Figure 2: Structure of Experiment: Instagram and Maps

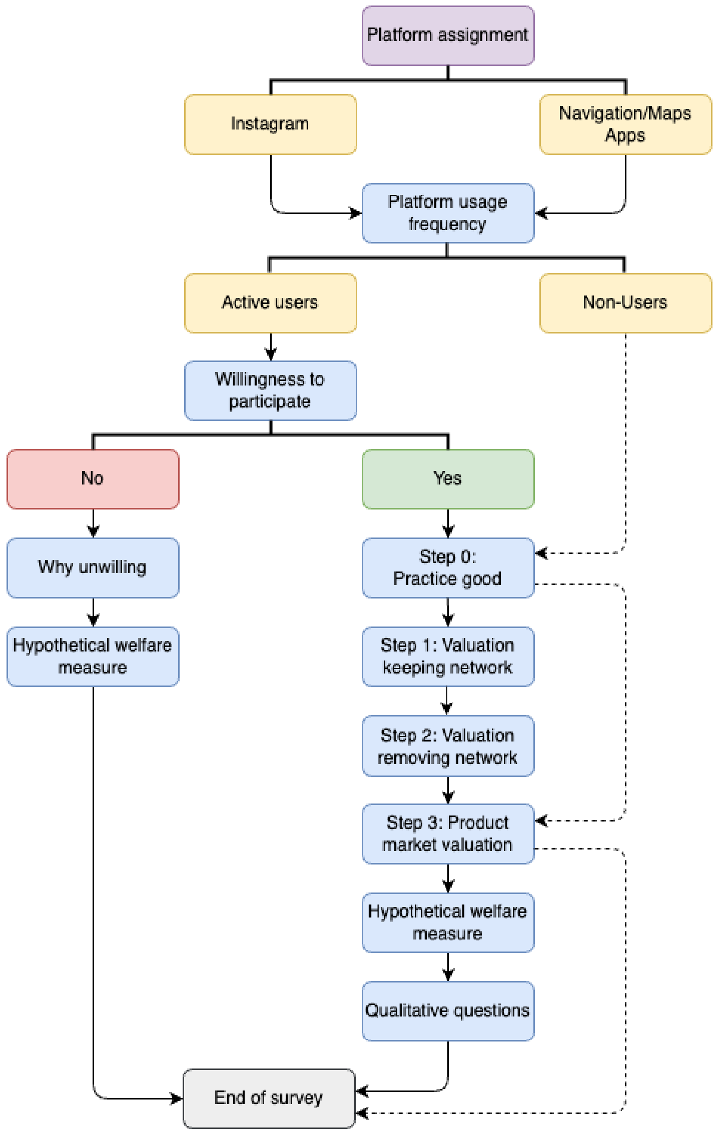

Notes: Figure 2 presents the structure of the experiment. At the beginning of the experiment, the platform is cross-randomized between Instagram and Maps. Active users and non-active users are directed to a distinct path. Active users are asked whether they are willing to participate in a deactivation study. The experiment ends for those unwilling to participate after two subsequent questions. The active users willing to participate are directed to steps 0 to 4, followed by the hypothetical welfare measure and a series of qualitative questions. Non-users proceed to steps 0 and 3, as indicated by the dashed arrows. The yellow boxes indicate embedded data, the blue boxes indicate question blocks, and the purple box indicates randomization. The flow of the TikTok experiment from July 2023 is identical except that there was no initial random platform assignment and that we did not elicit hypothetical welfare measures among respondents unwilling to participate in the study.

an iterative multiple price list, following the methodologies proposed by Andersen et al. (2006), and other best practices suggested in the literature (Berkouwer and Dean, 2022; Burchardi et al., 2021).

Our MPL places participants' valuation in one of 12 ranges, with lower and upper limits at \$0 and \$200 and internal increments of \$20: $(-\infty, \$0]$, $[\$0, \$20], \ldots, [\$180, \$200]$, $[\$200, \infty)$. In Step 3, we expand the limits to -\$200 and \$200, to account for the possibility of having a WTP as well as a WTA, resulting in 22 ranges. The algorithm proceeds sequentially, starting from an initial monetary offer and upper and lower bounds for the valuation. In each step, we present respondents with two options: in essence, either deactivation of their social media account and receiving the monetary offer, or keeping their social media account active. If the respondent accepts the offer (i.e., chooses to deactivate), her upper bound is set to that amount. Similarly, if she rejects the offer (i.e., keeps her account active), her lower bound is set to that amount. The algorithm then selects the next offer as the midpoint between her new bounds, resulting in progressively narrower valuation ranges with each response. The elicitation ends once we can narrow down the respondent's WTA to a \$20 range or once we surpass one of the upper or lower limits, which can take between 1 and 6 choices depending on the initial random offer and the respondent's choices.

To ensure that choices are incentive-compatible, we truthfully inform respondents that a computer will generate an amount of money to offer them to participate in the deactivation study. We further tell them that we will ask them a series of questions offering them different payment scenarios in case they are selected for the deactivation study. If they accept any price scenario lower than the computer's offer, we will invite them to the deactivation study and give them the computer's offer. If, on the other hand, they do not accept any price scenario lower than the computer's offer, we will not invite them to the deactivation study even if they are the selected participant. To examine comprehension, we ask respondents how demanding a higher amount affects their likelihood of receiving any payment. Reassuringly, 88% of respondents pass this comprehension check.

Step 0: Practice Good. To enhance comprehension, we start with a hypothetical example good (Dizon-Ross and Jayachandran, 2022). We measure individual-level willingness to accept the deactivation of respondents' ride-sharing accounts, taking aggregate consumption as given. We give them an initial offer randomized between \$0 and \$200 in \$20 increments and then have them decide between either (i) deactivating their ride-sharing accounts and sequentially varying amounts of money or (ii) not deactivating their accounts.

Step 1: Valuation Keeping Network. In step 1, we measure individuals' WTA to deactivate their social media accounts, taking aggregate consumption as given. We tell respondents that, to establish appropriate payment amounts for the deactivation study, we will ask them to decide whether to deactivate their social media account in exchange for different monetary amounts. We also reiterate that one student from their university will be randomly selected to participate in the study. We start the MPL with a randomly drawn offer between \$0 and \$200 in \$20 increments. Respondents then proceed to the multiple price list procedure where they decide between either (i) deactivating their social media account (with none of the students at their university deactivating) and sequentially varying amounts of money or (ii) not deactivating their account.[23]

Step 2: Valuation Removing Network. To assess the role of network effects in shaping individual consumer surplus, we measure individuals' valuation of their social media accounts when all participating students at their university are asked to deactivate their social accounts.

We start by truthfully presenting our participants with the possibility of a large-scale deactivation study at their university, where all participating students are asked to deactivate their accounts. In particular, we tell our respondents:

> College Pulse has a panel exceeding 650,000 university students. We are targeting universities with a high penetration of College Pulse.
>
> We will now ask you to consider two additional options for a large-scale deactivation of TikTok [Instagram] at your university. One of them will be randomly implemented if we manage to recruit more than two-thirds of the students at your university.
>
> We expect 90% of students to comply with deactivation based on previous studies (e.g., by Mosquera et al., 2018 and Allcott et al., 2020).

Thereafter, we inform respondents that we will randomly choose one of two options for conducting this deactivation study at a larger scale. We then proceed with describing

[23]To enhance comprehension and make the choices more intuitive, we added an explanation to the decisions of all 3 steps in terms of "taking a break from social media" or "not taking a break from social media" in the August and September 2023 collection. The collection from August and September 2023 also emphasizes more saliently on the decision screen that this choice is for the scenario in which the respondent would not receive any monetary payment.

the first option. We tell respondents that we will ask all participating students at their university sequentially whether they would like to deactivate their accounts. We then measure respondents' WTA to deactivate their social media accounts, conditional on us having asked all participating students at their university to deactivate their accounts in exchange for monetary payment. Respondents choose between (i) deactivating their account (when all other participating students have been also asked to deactivate) and receiving varying amounts of money sequentially and (ii) keeping their account active. To economize time, we randomize the initial offer between the lower and upper bounds of the respondent's valuation from step 1 (unless respondents are at the lower or upper ends of the WTA interval, in which case we offer them again this bound).[24]

Step 3: Product Market Valuation. In step 3, we measure the Product Market Valuation by eliciting individuals' preferences over the deactivation of the social media accounts of all participating students, including themselves.

Respondents are truthfully told that we know how much we need to pay every participating student at their university to deactivate their accounts for four weeks.[25] We inform respondents that we will randomly select one of the students to anonymously choose between the following two options: (i) keep things as they are or (ii) deactivate the accounts of all participating students. We clarify that if they decide for all participating students to deactivate their accounts, the researchers will pay the other students the amount they require. Moreover, they are told that we will establish their payment, if any, below.

To clarify the incentive compatibility of the mechanism, respondents learn that the deactivation study will be stopped for everyone only if they go back to using the platform before the end of the four weeks. In particular, the chosen respondent will not receive payment and we will pay the other students based on the actual time they spend in the study. Finally, if one of the other participating students goes back to using the platform before the end of the study, they will not receive any payment.

Subsequently, we remind people that the choice they make is incentivized. Respondents then proceed to the first main decision screen where they decide between (i) all participating students at their university deactivating their accounts (Option A) and (ii) all participating

[24]We randomize locally, as the step 1 valuation plausibly constitutes a more precise starting point for the step 2 elicitation compared to a fully randomized offer.

[25]Note that the information we collect through step 2 provides us with the necessary information to compensate respondents for their individual deactivation in the scenario of the large-scale deactivation that respondents decide upon in step 3. Since respondents in step 2 did not anticipate step 3, both elicitations are incentive-compatible.

students at their university keeping their accounts active (Option B). This first decision corresponds to the case in which the deciding participant does not receive payment, which effectively splits the participants' valuation into the positive or negative range. Consider the scenario where a respondent prefers Option A, of all participating students deactivating their accounts. On the following screen, they make a decision between all participating students deactivating their accounts vs. all participating students keeping their accounts active plus a random dollar amount, \$X, drawn between 0 and 200 in steps of 20. If the respondent chooses Option A once again, then she is willing to forgo a payment worth \$X. As in the previous steps, the subsequent offers are made iteratively to narrow down the respondent's WTP. Symmetrically, if she chooses Option B in the first screen, we then employ the iterative MPL algorithm to elicit her WTA to have all participating students deactivate their accounts.

Computing Welfare Based on the MPL. Responses to the MPL questions establish the lower and upper bounds of each respondent's WTA/WTP, effectively assigning them to one of the MPL ranges. For simplicity, we assign the mean of the endpoints for each range in order to have a unique WTA/WTP value; for instance, a range of [\$60, \$80] is assigned a value of \$70. The lowest and highest ranges are [-\$20, \$0] and [\$200, \$220], respectively.[26] In our robustness Section 3.4, we consider an alternative way of assigning participant's valuations.

3.3 The Social Media Trap

3.3.1 Main Results

We next proceed with presenting our main results for both TikTok and Instagram. We first proceed by presenting the traditional welfare measure not accounting for consumption spillovers to non-users. We then move on to the results of our new measure of welfare accounting for these spillovers. Finally, we present estimates of network effects.

Individual Consumer Surplus. The dark blue bars in Figure 3 show that the individual consumer surplus is large and positive, with a WTA to deactivate of approximately \$50 on average (with a median of \$30) for TikTok (Panel a) and Instagram (Panel b).[27] More-

[26]For Step 3, *Product Market Valuation*, the lowest range is expanded to [-\$220, -\$200].

[27]These estimates are lower than those in Allcott et al. (2020), who find a \$100 median valuation for Facebook. Aside from measuring welfare for different platforms, a possible explanation for the higher

over, as Figure 4 shows, roughly 90% of users derive positive welfare from both platforms, while the remaining respondents indicate requiring no payment for deactivating their account. This positive fraction of users requiring no payment could signal that some of them are partly aware of self-control problems and demand commitment devices, as documented by Allcott et al. (2022).

Product Market Surplus. We next turn to product market surplus, our preferred measure of welfare that accounts for consumption spillovers to non-users. The red bars in Figure 3 show that product market surplus is significantly lower compared to individual consumer surplus for both TikTok ($p < 0.01$) and Instagram ($p < 0.01$). Hence, the standard individual consumer surplus measure overstates welfare in this context, which is confirmed by Figure A4, where the inverse demand curve lies almost uniformly above the product market surplus curve. On average, users are *willing to pay* (rather than willing to accept) $28 and $10 to have others, including themselves, deactivate TikTok and Instagram, respectively. The median valuation is -$10 and $10 for TikTok and Instagram, respectively. As can be seen in Figure 4, 64% and 48% of active TikTok and Instagram users have a negative product market surplus. A similar pattern emerges for non-users: those without an account have an average WTP of $67 and $39 to have others deactivate TikTok and Instagram, respectively, with 87% and 56% of non-users of TikTok and Instagram exhibiting positive willingness to pay to have others deactivate their accounts.

The positive individual consumer surplus paired with the negative product market surplus is consistent with both users and non-users deriving, on average, a large and negative non-consumer surplus. In other words, many users report a large individual consumer surplus not because they derive a benefit from the platform, but because they would experience a negative utility if they were to be excluded from it. In that sense, a large fraction of active users are in a *social media trap*, as defined in our framework. Overall, these findings are evidence that the revealed-preference argument that users of a product derive positive welfare from it fails to apply in the presence of negative consumption spillovers to non-users. Our findings also highlight that college students are sophisticated about how others' social media consumption affects their own valuation.

valuation uncovered in Allcott et al. (2020) is that their sample consists of more active participants given their recruitment with Facebook ads, while we recruit college students.

Network Effects. Finally, we present our estimates of network effects, calculated by comparing the drop from the Valuation Keeping Network (Step 1), which reflects the individual consumer surplus, to the Valuation Removing Network (Step 2), which reflects the utility from using the product when the other participants are not using it.[28] Comparing the dark and light blue bars in Figure 3 uncovers a significant drop in valuation of 33% ($p < 0.01$) and 24% ($p < 0.01$) for TikTok and Instagram from Step 1 to Step 2, indicative of strong positive network effects. Moreover, and strikingly, only approximately 70% of users on both platforms derive positive welfare from the product in Step 2. This drop provides evidence that network effects are positive and quantitatively important, consistent with canonical theoretical frameworks (Rohlfs, 1974; Katz and Shapiro, 1985).

Our estimates also reveal that participants' average utility from using these platforms is positive when they are one of the few ones in their college using it, but negative in the status quo case where the rest of their school uses it as well. These results could be partly driven by our participants having a preference for the status of being one of the selected few ones in their school with access to the social media world. In this setting, aggregate use is not really zero (since the rest of the world keeps using the platform) and people might enjoy being the "gatekeepers" with access to the new trends when part of their network is excluded from the platform.

3.3.2 Correlates of Consumer Surplus

Table A3 examines heterogeneity in our different surplus measures along several demographics and displays regression coefficients from multivariate regressions. There are no significant correlations between age and gender and any of the surplus measures for both TikTok and Instagram.[29] Frequency in platform usage is positively significantly correlated with individual welfare measures for TikTok. Indeed, using the platform daily, as opposed to less frequently, is associated with a $24 increase in respondent's valuation for TikTok.

[28] Step 1 measures the marginal utility of using the platform given aggregate use X. Step 2 measures the marginal utility with lower aggregate use, so the difference between them reflects the definition of network effects that we proposed in Section 2.

[29] We assess the representativeness of our sample on these observables against the American Trends Panel of the Pew Research Center (2021). In the ATP data, 65% of TikTok users and 60% of Instagram users identify as female. In our sample, 74% of TikTok users and 66% of Instagram users identify as female; however, the percentage increases to 78% and 70% post attention and regret filters for TikTok and Instagram, respectively. Turning to age, 40% and 32% of TikTok and Instagram users, respectively, are aged between 18 and 29 in the ATP data. Naturally, this ratio is much higher in our sample of college students, where 97% and 96% of TikTok and Instagram users, respectively, are in this age group.

Figure 3: Consumer Surplus across Welfare Measures

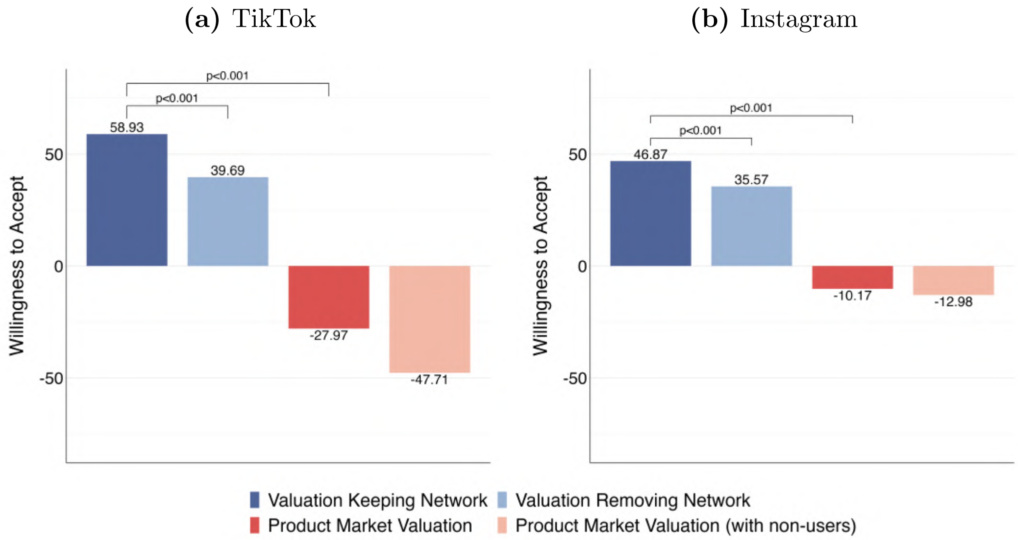

Notes: Figure 3 presents average valuations for the different welfare measures. Panel (a) presents the results for TikTok and Panel (b) presents the results for Instagram. The first three bars in each panel represent valuations exclusively for active users. The dark blue bar denotes *Valuation Keeping Network*; the light blue bar denotes *Valuation Removing Network*; the red bar denotes *Product Market Valuation* for users. The pink bar represents the average *Product Market Valuation* of active users and non-users. Respondents who agree with their elicited valuations and those who pass all of the attention checks are included. Reported p-values correspond to one-sided t-tests testing the null hypothesis that individual welfare estimates are lower than the aggregate welfare estimate.

Coefficients of daily platform usage are lower and more noisily measured for the product market valuation, compared to the individual measures.

3.4 Robustness

3.4.1 Measurement Differences Across Steps

Our *Valuation Keeping Network* does not allow for negative welfare, as it would require people being willing to pay for the deactivation of their account individually. As a result, there is naturally an asymmetry in measurement between *Valuation Keeping Network*, where those requiring no payment for deactivation are coded as having a valuation of -$10,[30] and the *Product Market Valuation*, where negative values of up to -$210 are possible.

To provide a very conservative way to examine whether this asymmetry in measurement can explain the sharp differences in valuation across *Valuation Keeping Network* and

[30]Reassuringly, only a small fraction of respondents (7% and 14% for TikTok and Instagram, respectively) require no payment for the deactivation of their account in *Valuation Keeping Network*.

Figure 4: Fraction with Negative Welfare across Welfare Measures

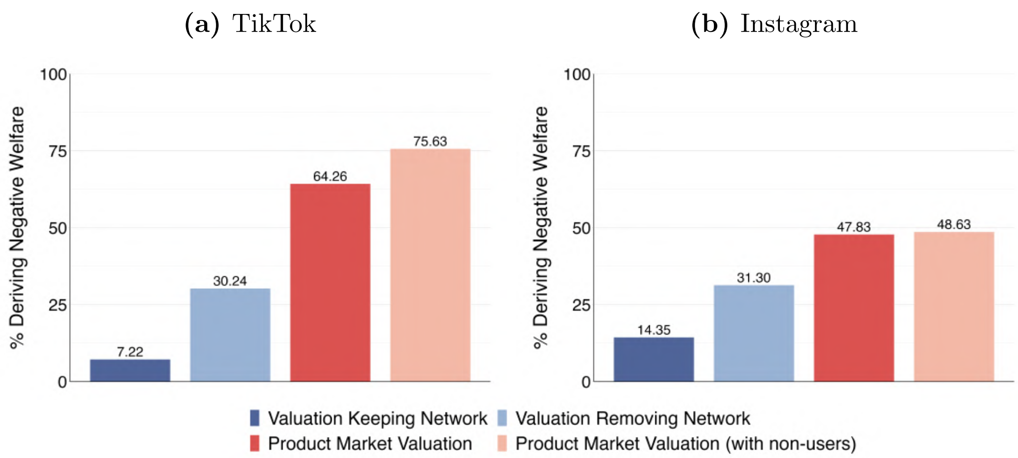

Notes: Figure 4 presents the percent of respondents with negative product valuations across our different welfare measures. Panel (a) presents the results for TikTok and Panel (b) presents the results for Instagram. The first three bars in each panel represent valuations exclusively for active users. The dark blue bar denotes *Valuation Keeping Network*; the light blue bar denotes *Valuation Removing Network*; the red bar denotes *Product Market Valuation* for users. The pink bar represents the average *Product Market Valuation* of active users and non-users. Respondents who agree with their elicited valuations and those who pass all of the attention checks are included.

Product Market Valuation, we conduct a simple bounding exercise. In that bounding exercise, we assume that participants that require no payment for the deactivation of their account in *Valuation Keeping Network* have a valuation of -\$210. Reassuringly, even under this very conservative bounding exercise, *Valuation Keeping Network* remains positive and large at \$45 for TikTok ($p < 0.05$) and \$18 for Instagram ($p < 0.05$), respectively.

3.4.2 Regret

To ensure data quality, we ask respondents whether they agree with a statement about what their choices mean in terms of their preferences over social media accounts for each of the four steps. For example, in the case of the practice good, a respondent with an implied valuation of between \$X1 and \$X2 is asked whether they agree with the statement that "According to your answers to the previous questions, you would require a payment worth between \$X1 and \$X2 to deactivate your Uber account for four weeks." If respondents do not agree with this statement, they are asked to complete the multiple price list questions one more time. Prior to the redirection, we inform participants that this will be their last chance to modify their answers. Our main sample is restricted to respondents who do not regret their final answers in any of the steps, but we discuss below that results still hold when we also include those that regret their final choice.

Overall, we find that 32% of respondents regret their choices once and a smaller fraction of 18% regret their choices twice. Figure A8 illustrates that this pattern holds for each step: after being redirected to the MPL questions, fewer participants disagree with their elicited WTA. The extent of regret fluctuates across steps. The percent of respondents regretting their choices is relative high at the practice section with 24% disagreeing with their elicited WTA initially and 15% after completing the MPL questions a second time. In the subsequent steps, the fraction of respondents regretting their final answers fluctuates around 10%. These patterns suggest that comprehension and data quality is high and that the practice questions helped improve comprehension.

3.4.3 Perceived Stakes and Credibility

One possible concern with our empirical design is that respondents may not find it likely that we will manage to recruit two-thirds of university students at their university, the condition for the large-scale deactivation study. To examine whether people perceived as credible that the large-scale deactivation study would take place, we asked respondents in our Instagram and Maps experiments about the percent chance that the researchers will recruit more than two-thirds of the students at their university. On average, participants perceive this likelihood to be quite substantial, at 45% for Instagram.[31] This in turn implies that respondents perceived the likelihood that the large-scale deactivation study would take place as substantial. Moreover, we examine how this perceived probability is correlated with our estimated welfare effects. Panel (e) of Figure A7 illustrates that individuals who deem the large-scale deactivation study as more probable have a more negative product market valuation. This heterogeneity suggests that our design is conservative and likely underestimates the extent of negative welfare. More broadly, given that even in the case of the large-scale deactivation study, not all users would deactivate their account, our study plausibly identifies lower bounds for the size of negative product market surplus and for the size of network effects.

3.4.4 Hypothetical Welfare Measures

Live in a World without. After the price elicitation, we present our respondents with a series of hypothetical qualitative questions. To assess the boundary conditions of our

[31] The perceived likelihood for Maps is reassuringly similar at 47%.

results (i.e., extrapolating to a hypothetical case where every user in the world stops using their social media), we ask respondents whether they would prefer to live in a world with or without the social media platform. As Figure 5 shows, 58% of respondents (including users and non-users) prefer to live in a world without TikTok and Instagram. Even among users, 33% and 57% prefer to live in a world without TikTok and Instagram, respectively.[32] Appendix Figure A7 validates these hypothetical survey questions with the incentivized measure of product market surplus. The figure illustrates that the hypothetical question is strongly correlated with the incentivized measure for both TikTok ($p < 0.01$) and Instagram ($p < 0.01$).

Figure 5: Percentage of Respondents that Prefer to Live in a World without the Platform

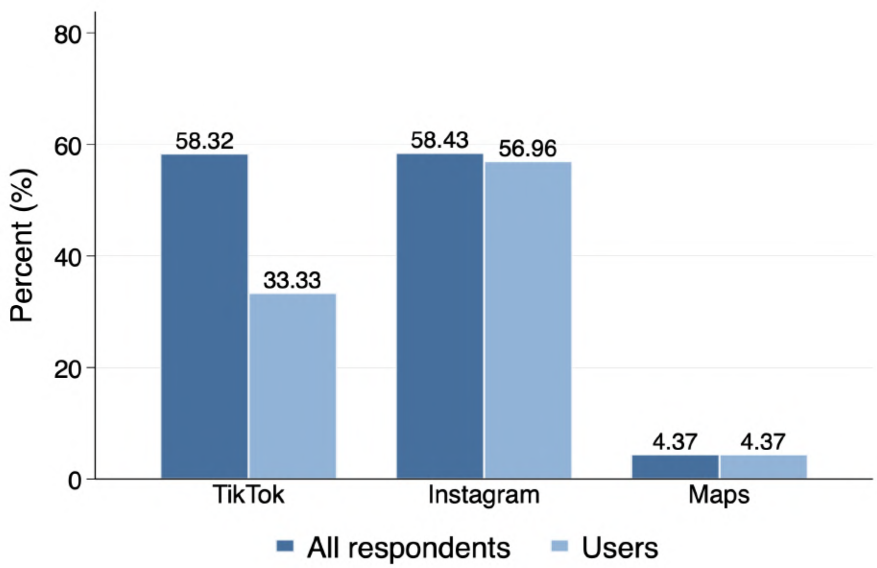

Notes: Figure 5 displays the percent of the respondents that stated they would prefer to live in a world without the platform for TikTok, Instagram and Maps separately. The dark blue bar represents the fraction among all respondents and the light blue bar represents the fraction among active users of the respective platform. Respondents who agree with their elicited valuations and those who pass all of the attention checks are included.

Preference Rankings. To understand the preferences of university students regarding social media platform usage among their peers, we ask them to rank three hypothetical

[32] As suggestive evidence against experimenter demand effects, Figure 5 also shows that only less than 5% of respondents prefer to live in a world without Maps.

scenarios: (i) they deactivate the platform and every other student at their university keeps using it, (ii) every student at their university, including themselves, deactivates the platform and (iii) no one deactivates the platform.

The results based on these rankings support our main findings. The most preferred scenario among our respondents is the scenario where every student at their universities, including themselves, deactivates their social media account, respectively (Figure A6). Among TikTok users, 39% prefer this option, while 49% of Instagram users prefer this option. In contrast, the least preferred scenario is where no one deactivates their account, with 49% and 52% of respondents citing this option as their least preferred one for TikTok and Instagram, respectively.[33]

Panels (c) and (d) of Appendix Figure A7 validate these hypothetical measures with the incentivized measure of product market surplus. This figure illustrates that the hypothetical question is strongly correlated with the incentivized measure. Indeed, while respondents who preferred deactivation for everyone have highly negative product market valuation for both Instagram and TikTok, respondents for whom deactivation for everyone was the least preferred option have positive product market valuation. The differences in product market surplus across these survey measures are highly significant for both Instagram and TikTok ($p < 0.01$).

3.4.5 Substitution Across Social Media Platforms

One concern is that people's product market valuation of social media platforms is so low given their opportunity to substitute their social media consumption to another platform, based on an argument similar to the one in Farrell and Saloner (1985). Specifically, even in the absence of non-consumer surplus, if there are two technologies, where an "alternative" technology is superior to the predominant technology, individuals' welfare could be improved if they all stopped using the inferior technology. Several pieces of evidence can help rule out this story and shed light on how substitution affects our estimates.

First, structural estimates of diversion ratios suggest that the outside option (offline or other online activities) is the most important substitution channel for social media, including Instagram and TikTok (Aridor, 2022). This pattern suggests that individuals do not all substitute towards a "better" platform. Second, as has been documented in the

[33] We also find that "Only I deactivate" is the preferred option for 35% and 30% of TikTok and Instagram users, respectively. This is consistent with our finding that a small fraction of users would accept to deactivate their accounts without compensation and with previous evidence on self-control problems.

literature (Allcott et al., 2022; Aridor, 2022), the vast majority of users in our data are multi-homers, with very few users having only a Tiktok (n=5 in the first survey) or an Instagram account (n=18 in the second survey). This adoption pattern makes it unlikely that users are trapped in Instagram or Tiktok because they cannot switch to a better alternative. Third, our estimates show that both respondents who multi-home and those that single-home have a negative product market valuation, which alleviates concerns that the negative product market surplus is driven by cross-platform substitution.[34]

3.4.6 Other Robustness Checks

Distributional Assumptions. As a robustness check against potential censoring in valuations, we assume a triangular distribution for those values that lie in these ranges, following the methodology of Allcott and Kessler (2019). Estimates constructed this way give more weight to the upper and lower bounds in the elicitation and thus allow us to gauge the sensitivity of our results to extreme valuations at the tails. Given that we see more mass at the lower end of the distribution (see Appendix Figure A5) this means that the welfare estimates based on the triangular distribution are more negative than our main estimates (see Appendix Figure A13). This, in turn suggests, that censoring, if anything, makes us overestimate the welfare effects of social media. See Appendix B.1.2 for details on the triangular distribution.

Robustness to Sample Restrictions. In our main analysis, we reported results for respondents who passed all attention checks and did not regret any of their final choices. Appendix Figure A10 confirms our results for the full sample without those exclusions; Appendix Figure A11 confirms our findings with a sample that includes inattentive respondents and excludes respondents regretting their final choice. Finally, Appendix Figure A12 demonstrates the robustness of our findings to including respondents regretting their final choice and only excluding inattentive respondents.

[34]Among individuals that have only a TikTok account (n=5) and only an Instagram account (n=18), the product market surplus is even more negative at -$54 and -$57, respectively. Naturally, these estimates are noisy given that most respondents have both a TikTok and an Instagram account. Among users that multi-home, the estimates are -$28 for TikTok and -$6 for Instagram.

3.5 Mechanisms

3.5.1 Repugnance Towards Digital Products

One possible mechanism that might explain the drop in welfare from the individual to the product market surplus could be repugnance towards digital products. To test this conjecture, we run a deactivation study experiment with a digital good that plausibly does not cause strong consumption spillovers, Maps; these likely have more muted consumption spillovers on non-users as they do not create social costs of exclusion and are less likely to impact relative social standing. The instructions are virtually identical to our main experiment, except for the product name.

Figure A9 shows that both the individual consumer surplus and the product market surplus are positive and significantly different from zero in the case of Maps.[35] The product market surplus is lower than the individual consumer surplus, which might result from various motives. First, respondents may dislike "Big Tech" companies and their associated market power and therefore prefer a ban of products that underlie the market power of big tech companies. Second, respondents may feel repugnance towards digital products, such as mobile phones and modern technologies. Third, respondents may have a distaste for others using their phone.

To formally test whether the drop in welfare between individual consumer surplus and product market surplus is larger for Instagram compared to Maps, we conduct a simple difference-in-differences exercise, where we compute the change between valuation keeping network and product market valuation between Instagram and Maps. Table A2 reports estimates of the difference-in-differences regression and shows that the coefficient on the interaction term (of an indicator of Instagram relative to Maps and product) is of substantial magnitude and highly significant. This corroborates that negative consumption spillovers for non-users are larger on Instagram than they are for Maps.

The evidence on the positive product market valuation of Maps and the significant difference-in-differences estimate also alleviates concerns that the question wording we used in the product market valuation mechanically induces negative welfare estimates, i.e. respondents providing positive willingness to pay to ban others using the product.

[35]These findings are also consistent with the hypothetical ranking question, where we find that the most preferred scenario is for no one to quit maps (46%), while only 24% of respondents have a preference for everyone to quit maps (see Appendix Figure A9d).

3.5.2 Motives Behind Consumption

Social Media Platforms. We next turn to more direct evidence on mechanisms. In particular, we asked active users of the platform who said that they prefer to live in a world without the platform an open-ended question to better understand the motives behind their usage.[36] We asked them the following question:

> You mentioned you would prefer to live in a world without [platform]. Why do you still use it?

To quantitatively analyze the data, we devised a simple hand-coding scheme, which comprises five categories.[37] "FOMO" responses usually mention feeling left out ("I feel like if I stop using it, I will be completely out of the loop").[38] "Entertainment" responses talk about the high entertainment value of the platform ("It's a very good source of entertainment and it's always something to do when bored"). "Addiction" responses mention self-control problems and addiction ("It's very addicting and I cannot stop"). "Information" responses indicate receiving useful information ("I follow pages that keep me up to date with the largest news"). Finally, "Productivity/Convenience" responses mention using the platform for productive use or convenience ("I still use Instagram for business purposes"). Appendix Table A4 provides an overview of the hand-coding scheme and provides further example responses.[39]

Figure 6 illustrates the quantitative distribution of the hand-coded data for TikTok, Instagram and Maps, respectively. It reveals that the fear of missing out is the most prevalent motive both for Instagram (79%) and TikTok (40%). Moreover, entertainment motives also play an important role in driving people's social media consumption (36% of TikTok and 21% of Instagram users), consistent with evidence on people's news preferences

[36] Open-ended questions are increasingly used to better understand the hidden motives behind people's choices, see e.g. Bursztyn et al. (2023b, 2022). These questions avoid priming respondents on particular motivations and better capture what naturally comes to mind compared to more structured questions (Haaland et al., 2023).

[37] A given response can fall into multiple categories.

[38] Social media platforms like TikTok and Instagram may cultivate FOMO among users through specific platform features and content dynamics. TikTok's video format and evolving trends may create social pressure to remain continually informed. Conversely, Instagram's feature of shared content that is only available for a limited time on the platform may foster a fear of being left out.

[39] We validate our hand-coded open-ended data with data coded by a large language model (LLM). We show that the LLM-based measure yields similar frequencies of the categories (See Panel (a) of Figure A14). Moreover, the LLM-based categories are highly correlated with the hand-coded measure (see Panel (a) of Table A6 for details).

(Bursztyn et al., 2023a). Consistent with prior evidence (Allcott et al., 2022), addiction is an important reason for TikTok (35%), though somewhat less important for Instagram (10%). Obtaining information is also an important reason, with 16% of TikTok and 15% of Instagram users citing this as the reason. Finally, only a very small fraction of users (3% and 7% on TikTok and Instagram, respectively) cite productivity/convenience as a reason for using the platform.

Maps. We also conducted a similar coding procedure for the Maps experiment. Figure 6 reveals a very different distribution of motives: 60% of respondents mention productivity reasons, 30% mention information and only 10% mention the fear of missing out ("[...] still use navigation maps because it is what everyone uses"). The strikingly different patterns for navigation apps compared to social media apps are suggestive evidence against experimenter demand effects driving the uncovered patterns (de Quidt et al., 2018).

3.5.3 Direct Evidence on Mechanisms Behind Consumption Spillovers

Social Media Platforms. To provide direct evidence on the mechanisms behind consumption spillovers to non-users, we asked all of our respondents an open-ended question to describe the nature and motives behind their non-user surplus. In particular, we asked them "how they would feel if they were the only one who deactivated [platform] and everyone else kept using it?."

Based on the open-ended responses we devised a coding scheme to capture the most common topics. "FOMO" responses talk about the fear of missing out ("I would definitely feel a bit left out."). "Negative" responses express negative emotions without explicitly mentioning the fear of missing out ("[...] it would be a little unfair"). "Indifferent" responses indicate that they do not expect the deactivation to have strong effects on them ("That wouldn't be a big deal"). "Beneficial" responses mention the benefits of not using the respective platforms ("I would be able to focus on more important things").[40]

Figure 7 illustrates the results. Panel (a) shows results for TikTok among respondents who prefer to live in a world without TikTok.[41] Panel (b) shows results for Instagram, separately for respondents who prefer to live in a world with and without Instagram. Among the TikTok users who would prefer to live in a world without TikTok, 36% express

[40] We again show that the hand-coded measure is highly correlated with analogous data annotated by a large language model (see Panel (b) of Table A6 and Panel (b) of Figure A14.)

[41] We did not collect the open-ended data for respondents who preferred to live in a world with TikTok.

Figure 6: Motives for Social Media Consumption Despite a Preference to Live in a World without It

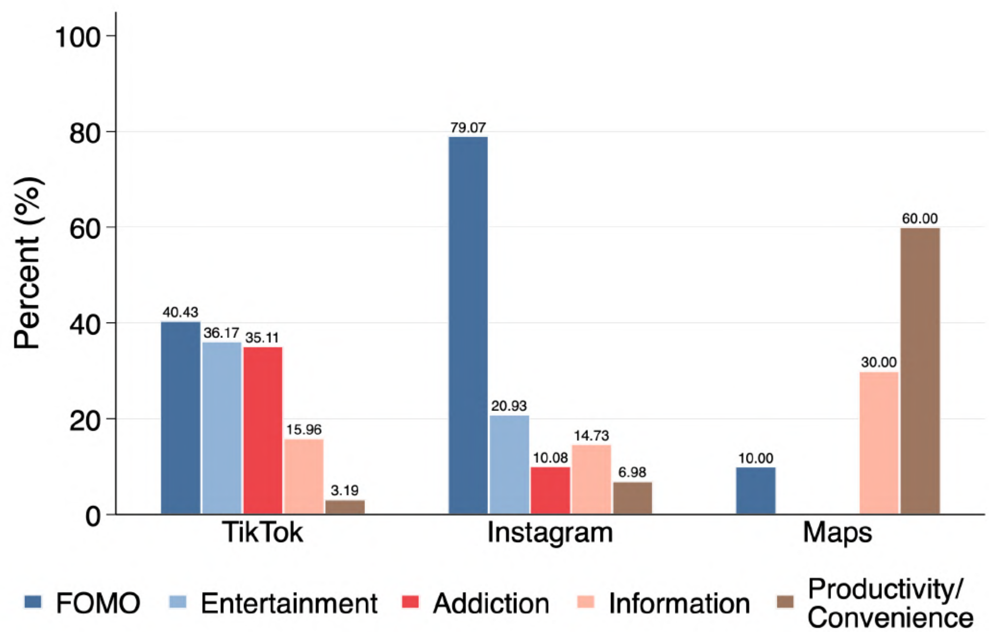

Notes: Figure 6 presents the fraction of respondents mentioning different motives in their open-ended responses. Active users who said that they prefer to live in a world without the platform were asked the following open-ended question: *You mentioned you would prefer to live in a world without [platform]. Why do you still use it?* "FOMO" denotes responses mentioning the fear of missing out or related social concerns. "Entertainment" denotes responses mentioning the entertainment value of the platform. "Addiction" denotes responses indicating the addictive nature of the platform and self-control problems. "Information" denotes responses mentioning informational purposes such as following the news or keeping abreast of college events. "Productivity" denotes responses mentioning productivity benefits, such as using the platform for business purposes. The categorization of the open-ended answers is not mutually exclusive. Respondents who agree with their elicited valuations and those who pass all of the attention checks are included. Nonsensical responses were dropped from the analysis. The underlying sample sizes are 94 for TikTok, 129 for Instagram, and 10 for Maps.

FOMO, 29% are indifferent, 9% have generic negative feelings and a 21% see it as beneficial. Among the Instagram users who would prefer to live in a world without Instagram, 38% express FOMO, 34% are indifferent, 8% have generic negative feelings, and a 20% see it as beneficial. Next, we turn to Instagram users who would prefer to live in a world with Instagram. Even among those, 40% express FOMO, 26% are indifferent, 14% have generic negative feelings, and 16% see it as beneficial. These data clearly highlight that the fear of missing out is a prevalent motive behind non-users' consumption spillovers across people who prefer to live in a world with and without Instagram.

Maps. We hypothesized that the nature of consumption spillovers for navigation and maps smartphone applications is different from social media platforms. To provide direct evidence for this conjecture, we also analyzed people's responses to the open-ended question for Maps. Figure 7 includes patterns separately for respondents who prefer to live in a world with and without maps. Among respondents preferring to live in a world without Maps, FOMO is mentioned fairly infrequently (10%)[42], while "Indifferent" responses and "Negative" responses are more prevalent at 30 and 40%, and 20% of responses fall into the "Beneficial" category. Among respondents preferring to live in a world with Maps, the patterns are similar: 15% mention FOMO, 27% are "Indifferent" responses, 44% are "Negative" responses, and 9% are "Beneficial" responses. Overall, these patterns underscore that the mechanisms behind consumption spillovers to non-users are very different for social media apps and navigation apps and that social forces, like the fear of missing out, play a much more dominant role in the context of social media.

4 Other Applications

The evidence on product market traps in the previous section is specific to social media and mechanisms related to the fear of missing out. To probe how common these forces are, we provide suggestive survey evidence on the existence of product market traps for luxury and preferences about the frequency of product variations.

[42] An example response is: "would feel a bit isolated, maybe excluded from certain conversations involving travel plans, etc."

Figure 7: Evidence on Mechanisms Behind Consumption Spillovers to Non-Users

(a) Active Users that Prefer to Live in a World without Platform

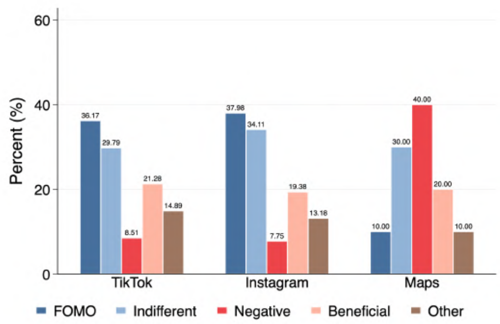

(b) Active Users that Prefer to Live in a World with Platform

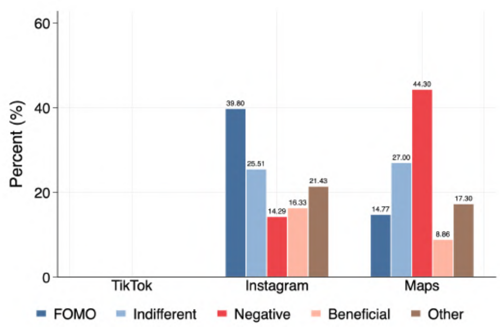

Notes: Figure 7 presents the fraction of respondents expressing different emotions in their open-ended responses. Panel (a) shows results for respondents who prefer to live in a world without the respective platform, while Panel (b) shows results for those who prefer to live in a world with the respective platform. Active users were asked the following open-ended question: *How would you feel if you were the only one who deactivated [platform] and everyone else kept using it?* Data for TikTok is missing for Panel (b) as this question was only directed to TikTok users who stated they would rather live in a world without TikTok. "FOMO" denotes responses mentioning the fear of missing out or related social concerns. "Indifferent" denotes responses expressing they would not be particularly affected. "Negative" denotes responses expressing negative emotions, whereas "Beneficial" denotes responses where respondents mention a potential benefit of deactivation. "Other" denotes a diverse set of responses that mention different motives. The categorization of the open-ended answers is not mutually exclusive. Respondents who agree with their elicited valuations and those who pass all of the attention checks are included. Nonsensical responses were dropped from the analysis. Responses indicating indifference conditional on payment/contribution to research were placed in the "Other" category. The underlying sample sizes are 94 for TikTok, 227 for Instagram, and 247 for Maps.

4.1 Luxury Goods as Product Market Traps

We start with evidence on luxury goods, where positional externalities are a plausible driver of consumption spillovers to potential non-consumers.

Sample. We fielded pre-registered surveys with 500 US participants from Prolific, a widely used online labor market used for social science experiments (Eyal et al., 2021), in September 2023.[43]

Survey. Our survey consists of two blocks: one block on luxury goods discussed in this section and another block on vintage goods presented in Section 4.2.[44] In the luxury block, we ask respondents to indicate whether they owned products from luxury brands they personally purchased.[45] We then leverage the hypothetical questions that we validated with incentivized measures of welfare in our social media experiments. In particular, we ask respondents whether they prefer to live in a world with or without any of these luxury brands. The full set of instructions can be found in Appendix C.3.

Results. In our survey, 34% of respondents own luxury brands. Conditional on owning any luxury brand, they owned 2.1 luxury brands on average. Figure 8 shows that among respondents who owned any goods of luxury brands, 44% preferred to live in a world without those brands. Among respondents not owning any of these brands, the fraction preferring to live in a world without them is higher, at 69%. While the literature on luxury goods has emphasized the negative externalities these goods impose on non-consumers (Frank, 1985a, 2000, 2012), our evidence highlights that large shares of consumers of these products would prefer them not to exist. Given that these results are in line with our social media estimates, it is plausible that status concerns might be an important mechanism driving negative non-consumer surplus.

4.2 Frequency of Product Variations

Product market traps lead to a situation where the existence of a product is harmful to consumers. This can manifest as excessive consumption by users or the production of an

[43]The pre-registration can be found on AsPredicted #144630.

[44]We randomize the order between these two blocks.

[45]The brands we used are: Louis Vuitton, Gucci, Chanel, Yves Saint Laurent (YSL), Balenciaga, Versace, Rolex, Tiffany & Co., Burberry, Givenchy, and Swarovski. Respondents could also fill in any other luxury brand not part of this list.

Figure 8: Luxury and Vintage Goods

(a) Percentage of Respondents that Prefer to Live in a World Without Any Luxury Brands

(b) Percentage of Respondents that Prefer to Live in a World Where iPhone is Released Every Other Year

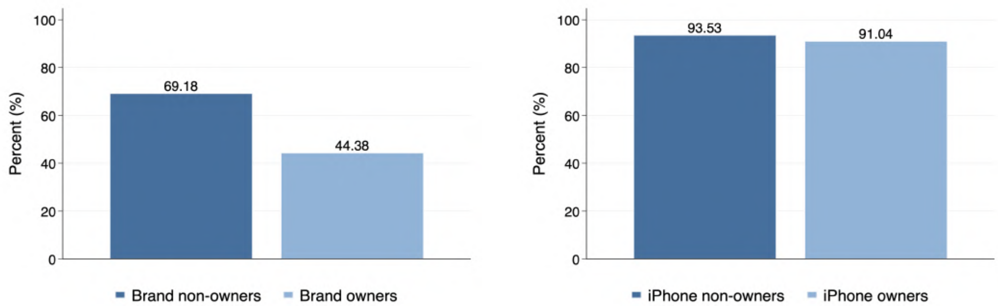

Notes: Panel (a) of Figure 8 displays the fraction of respondents preferring to live in a world without any luxury brands separately for brand owners and brand non-owners. Panel (b) displays the fraction of respondents that prefer to live in a world where Apple releases the new iPhone every other year rather than every year, separately for Phone owners and iPhone non-owners.

excessive number of product variations or vintages (Pesendorfer, 1995).

To examine people's preferences regarding the frequency of product variations, we asked respondents whether they would prefer to live in a world where Apple releases the iPhone every year or every other year in the survey presented in the previous section. We document that, among iPhone owners, a striking 91% of respondents would prefer Apple to release the iPhone every other year rather than every year.[46] Among respondents not owning the iPhone, 94% prefer Apple to release the iPhone every other year rather than every year. This finding provides suggestive evidence that consumers consider the number of product variations or vintages of the iPhone as excessive and thus harmful to consumer welfare. Overall, the findings from this survey suggest that negative non-consumer surplus is not specific to the case of social media, but also extends to luxury consumption and particular high-end technology products.

5 Conclusion

In the conventional assessment of consumer welfare, the emphasis is predominantly on individual-level evaluations, holding aggregate consumption fixed. However, our findings

[46]It is worth highlighting that among iPhone owners only 8% prefer to live in a world without iPhones, while among respondents not owning the iPhone this fraction is 49%.

show that such measures do not accurately reflect welfare in settings with consumption spillovers to potential non-users. Consequently, we introduce a new method to gauge welfare in these contexts, which we apply to widely used social media platforms through large-scale incentivized trials involving college students. While traditional measures of individual consumer surplus suggest positive welfare, the *Product Market Valuation* that factors in consumption spillovers to non-users tells a different story: it reveals negative welfare, with a notable portion of users experiencing a net disutility from the platform.

Our theoretical framework introduces the concept of *product market traps*, a phenomenon whereby consumers find themselves trapped in an inefficient market equilibrium and prefer the product not to exist, but cannot coordinate to stop using it. Intriguingly, such product market traps can arise even with fully rational expectations and without any behavioral frictions. In the context of social media, our empirical evidence highlights the existence of a *social media trap* for a large fraction of consumers, who derive large individual consumer surplus but, simultaneously, experience negative welfare from the product. These results could help reconcile the seemingly contradictory findings in the social media literature of a large consumer surplus coexisting with negative effects on well-being. More generally, these findings challenge the standard revealed-preference argument that the mere existence of a product implies that its consumers derive positive welfare. The presence of product market traps underscores the need for more research on whether companies introduce features that exacerbate non-consumer surplus and diminish consumer welfare, rather than enhance it, increasing people's need for a product without increasing the utility it delivers to them.

Our framework also highlights a few important levers for policymakers: first, policymakers should regulate markets to counteract producers' incentives to use technologies that decrease non-consumer surplus. Second, given that larger networks decrease non-consumer surplus, optimal anti-trust policy may involve reducing the size of networks.

References

Acemoglu, Daron, Ali Makhdoumi, Azarakhsh Malekian, and Asu Ozdaglar, "Too much data: Prices and inefficiencies in data markets," *American Economic Journal: Microeconomics*, 2022, *14* (4), 218–256.

Akerlof, Robert, Richard Holden, and Luis Rayo, "Network externalities and market dominance," *Management Science*, 2023.

Alevy, Jonathan E, John A List, and Wiktor L Adamowicz, "How can behavioral economics inform nonmarket valuation? An example from the preference reversal literature," *Land Economics*, 2011, *87* (3), 365–381.

Allcott, Hunt and Judd B Kessler, "The welfare effects of nudges: A case study of energy use social comparisons," *American Economic Journal: Applied Economics*, 2019, *11* (1), 236–276.

_ , Luca Braghieri, Sarah Eichmeyer, and Matthew Gentzkow, "The welfare effects of social media," *American Economic Review*, 2020, *110* (3), 629–676.

_ , Matthew Gentzkow, and Lena Song, "Digital addiction," *American Economic Review*, 2022, *112* (7), 2424–2463.

Ambuehl, Sandro, B Douglas Bernheim, and Axel Ockenfels, "What motivates paternalism? An experimental study," *American economic review*, 2021, *111* (3), 787–830.

American Trends Panel of the Pew Research Center, "Wave 93," 2021. Accessed: 2023-10-02.

Andersen, Steffen, Glenn W Harrison, Morten Igel Lau, and E Elisabet Rutström, "Elicitation using multiple price list formats," *Experimental Economics*, 2006, *9*, 383–405.

Aridor, Guy, "Drivers of digital attention: Evidence from a social media experiment," 2022.

Bagwell, Laurie Simon and B Douglas Bernheim, "Veblen effects in a theory of conspicuous consumption," *The American economic review*, 1996, pp. 349–373.

Becker, Gary S, "A note on restaurant pricing and other examples of social influences on price," *Journal of political economy*, 1991, *99* (5), 1109–1116.

Becker, Gordon M, Morris H DeGroot, and Jacob Marschak, "Measuring utility by a single-response sequential method," *Behavioral science*, 1964, *9* (3), 226–232.

Benzell, Seth and Avinash Collis, "Regulating Digital Platform Monopolies: The Case of Facebook," in "Regulating Digital Platform Monopolies: The Case of Facebook: Benzell, Seth— uCollis, Avinash," [Sl]: SSRN, 2022.

Berkouwer, Susanna B and Joshua T Dean, "Credit, attention, and externalities in the adoption of energy efficient technologies by low-income households," *American Economic Review*, 2022, *112* (10), 3291–3330.

Bernheim, B Douglas, "The good, the bad, and the ugly: A unified approach to behavioral welfare economics," *Journal of Benefit-Cost Analysis*, 2016, *7* (1), 12–68.

_ **and Dmitry Taubinsky**, "Behavioral Public Economics," *Handbook of Behavioral Economics: Applications and Foundations 1*, 2018, *1*, 381–516.

Bhattacharya, Debopam, Pascaline Dupas, and Kanaya Shin, "Demand and Welfare Analysis in Discrete Choice Models with Social Interactions," *The Review of Economic Studies*, 2023.

Bottan, Nicolas L and Ricardo Perez-Truglia, "Choosing your pond: location choices and relative income," *Review of Economics and Statistics*, 2022, *104* (5), 1010–1027.

Braghieri, Luca, Ro'ee Levy, and Alexey Makarin, "Social media and mental health," *American Economic Review*, 2022, *112* (11), 3660–3693.

Brynjolfsson, Erik and JooHee Oh, "The attention economy: Measuring the value of free digital services on the internet," *Information Systems Research*, 2023.

_ **, Avinash Collis, and Felix Eggers**, "Using massive online choice experiments to measure changes in well-being," *Proceedings of the National Academy of Sciences*, 2019, *116* (15), 7250–7255.

_ **,** _ **, Asad Liaqat, Daley Kutzman, Haritz Garro, Daniel Deisenroth, Nils Wernerfelt, and Jae-Joon Lee**, "The Digital Welfare of Nations: New Measures of Welfare Gains and Inequality," Working Paper 31670, National Bureau of Economic Research 2023.

Burchardi, Konrad B, Jonathan De Quidt, Selim Gulesci, Benedetta Lerva, and Stefano Tripodi, "Testing willingness to pay elicitation mechanisms in the field: Evidence from Uganda," *Journal of Development Economics*, 2021, *152*, 102701.

Bursztyn, Leonardo, Aakaash Rao, Christopher Roth, and David Yanagizawa-Drott, "Opinions as facts," *The Review of Economic Studies*, 2023, *90* (4), 1832–1864.

_ **and David Y Yang**, "Misperceptions about others," *Annual Review of Economics*, 2022, *14*, 425–452.

_ **, Bruno Ferman, Stefano Fiorin, Martin Kanz, and Gautam Rao**, "Status goods: experimental evidence from platinum credit cards," *The Quarterly Journal of Economics*, 2018, *133* (3), 1561–1595.

_ **, Florian Ederer, Bruno Ferman, and Noam Yuchtman**, "Understanding mechanisms underlying peer effects: Evidence from a field experiment on financial decisions," *Econometrica*, 2014, *82* (4), 1273–1301.

— , Georgy Egorov, Ingar Haaland, Aakaash Rao, and Christopher Roth, "Scapegoating during Crises," *AEA Papers and Proceedings*, 2022, *112*, 151–55.

— , — , — , — , and — , "Justifying dissent," *The Quarterly Journal of Economics*, 2023, *138* (3), 1403–1451.

Butera, Luigi, Robert Metcalfe, William Morrison, and Dmitry Taubinsky, "Measuring the welfare effects of shame and pride," *American Economic Review*, 2022, *112* (1), 122–168.

Cherye, Cristobal and Alessandro Acquisti, "Online intermediation in legacy industries: The effect of reservation platforms on restaurants' prices and survival," 2022.

Choi, Jay Pil, Doh-Shin Jeon, and Byung-Cheol Kim, "Privacy and personal data collection with information externalities," *Journal of Public Economics*, 2019, *173*, 113–124.

Clark, Andrew E and Andrew J Oswald, "Satisfaction and comparison income," *Journal of public economics*, 1996, *61* (3), 359–381.

Cullen, Z. and R. Perez-Truglia, "How Much Does Your Boss Make? The Effects of Salary Comparisons," *Journal of Political Economy*, 2022, *130* (3), 766–822.

de Quidt, Jonathan, Johannes Haushofer, and Christopher Roth, "Measuring and Bounding Experimenter Demand," *American Economic Review*, November 2018, *108* (11), 3266–3302.

DellaVigna, Stefano, John A List, and Ulrike Malmendier, "Testing for altruism and social pressure in charitable giving," *The quarterly journal of economics*, 2012, *127* (1), 1–56.

Dizon-Ross, Rebecca and Seema Jayachandran, "Improving Willingness-to-Pay Elicitation by Including a Benchmark Good," in "AEA Papers and Proceedings," Vol. 112 American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203 2022, pp. 551–555.

Eyal, Peer, Rothschild David, Gordon Andrew, Evernden Zak, and Damer Ekaterina, "Data quality of platforms and panels for online behavioral research," *Behavior Research Methods*, 2021, pp. 1–20.

Faelens, Lien, Kristof Hoorelbeke, Ruben Cambier, Jill van Put, Eowyn Van de Putte, Rudi De Raedt, and Ernst HW Koster, "The relationship between Instagram use and indicators of mental health: A systematic review," *Computers in Human Behavior Reports*, 2021, *4*, 100121.

Farrell, Joseph and Garth Saloner, "Standardization, compatibility, and innovation," *the RAND Journal of Economics*, 1985, pp. 70–83.

_ **and Paul Klemperer**, "Coordination and lock-in: Competition with switching costs and network effects," *Handbook of industrial organization*, 2007, *3*, 1967–2072.

Ferrara, La, E Gulesci S, D Smerdon, and M Sulaiman, "Changing harmful norms through information and coordination: Experimental evidence from somalia," 2023.

Frank, Robert H, *Choosing the right pond: Human behavior and the quest for status.*, Oxford University Press, 1985.

_ , "The demand for unobservable and other nonpositional goods," *The American Economic Review*, 1985, *75* (1), 101–116.

_ , *Luxury fever: Money and happiness in an era of excess*, Princeton University Press, 2000.

_ , "Positional externalities cause large and preventable welfare losses," *American economic review*, 2005, *95* (2), 137–141.

_ , "The Darwin Economy," in "The Darwin Economy," Princeton University Press, 2012.

Granovetter, Mark, "Threshold models of collective behavior," *American journal of sociology*, 1978, *83* (6), 1420–1443.

Gupta, Mayank and Aditya Sharma, "Fear of missing out: A brief overview of origin, theoretical underpinnings and relationship with mental health," *World journal of clinical cases*, 2021, *9* (19), 4881.

Haaland, Ingar, Christopher Roth, and Johannes Wohlfart, "Designing information provision experiments," *Journal of economic literature*, 2023, *61* (1), 3–40.

Heffetz, Ori and Robert H Frank, "Preferences for status: Evidence and economic implications," in "Handbook of social economics," Vol. 1, Elsevier, 2011, pp. 69–91.

Imas, Alex and Kristóf Madarász, "Superiority-Seeking and the Preference for Exclusion," *Review of Economic Studies*, 2022.

Jullien, Bruno, Alessandro Pavan, and Marc Rysman, "Two-sided markets, pricing, and network effects," in Kate Ho, Ali Hortacsu, and Alessandro Lizzeri, eds., *Handbook of Industrial Organization, Volume 4*, Vol. 4 of *Handbook of Industrial Organization*, Elsevier, 2021, pp. 485–592.

Katz, Michael L and Carl Shapiro, "Network externalities, competition, and compatibility," *The American economic review*, 1985, *75* (3), 424–440.

Kemp, Simon, "Digital 2022: Global Overview Report," Available at https://datareportal.com/reports/digital-2022-global-overview-report 2022.

Kuhn, Peter, Peter Kooreman, Adriaan Soetevent, and Arie Kapteyn, "The effects of lottery prizes on winners and their neighbors: Evidence from the Dutch postcode lottery," *American Economic Review*, 2011, *101* (5), 2226–2247.

Leibenstein, Harvey, "Bandwagon, snob, and Veblen effects in the theory of consumers' demand," *The Quarterly Journal of Economics*, 1950, *64* (2), 183–207.

Luttmer, Erzo FP, "Neighbors as negatives: Relative earnings and well-being," *The Quarterly journal of economics*, 2005, *120* (3), 963–1002.

Manski, Charles F, "Identification of endogenous social effects: The reflection problem," *The review of economic studies*, 1993, *60* (3), 531–542.

Mosquera, Roberto, Mofioluwasademi Odunowo, Trent McNamara, Xiongfei Guo, and Ragan Petrie, "The economic effects of Facebook," *Experimental Economics*, 2020, *23*, 575–602.

Perez-Truglia, Ricardo, "The effects of income transparency on well-being: Evidence from a natural experiment," *American Economic Review*, 2020, *110* (4), 1019–1054.

Pesendorfer, Wolfgang, "Design innovation and fashion cycles," *The american economic review*, 1995, pp. 771–792.

Pew Research Center, "Social Media Use in 2021," 2021. Accessed: 2023-08-22.

Rochet, Jean-Charles and Jean Tirole, "Platform competition in two-sided markets," *Journal of the european economic association*, 2003, *1* (4), 990–1029.

Rohlfs, Jeffrey, "A theory of interdependent demand for a communications service," *The Bell journal of economics and management science*, 1974, pp. 16–37.

Romer, Paul, "A Tax That Could Fix Big Tech.," *International New York Times*, 2019, pp. NA–NA.

Roth, Alvin E., "Repugnance as a Constraint on Markets," *Journal of Economic Perspectives*, September 2007, *21* (3), 37–58.

Sullivan, Michael, "Price controls in a multi-sided market," Technical Report, Working paper 2022.

Tirole, Jean, "Competition and the industrial challenge for the digital age," *Annual Review of Economics*, 2023, *15*, 573–605.

Online Appendix: Not for publication

Our supplementary material is structured as follows. Section A provides additional details related to the conceptual framework. Section B includes additional tables and figures. Section B.1 provides additional evidence on robustness. Finally, Appendix C presents the instructions for all experiments described in the paper.

A Additional framework details

Changes in the equilibrium use. Our model specified above allows for comparative statics on how welfare changes for distinct groups (e.g., users who lose out from the product) as equilibrium use X^* changes. Figure A1 provides an illustrative example of how an increase in the number of users can, simultaneously, 1) increase the marginal utility from using the product relative to not using it (positive network effects) and 2) decrease the utility of both using and not using the product. Our empirical application to social media elicits portions of these two curves using experimental methods, allowing us to assess key comparative statics related to how welfare changes with total equilibrium use.

Appendix Figure A1: Utility and Negative Consumption Spillovers from Product Use as a Function of Aggregate Use X

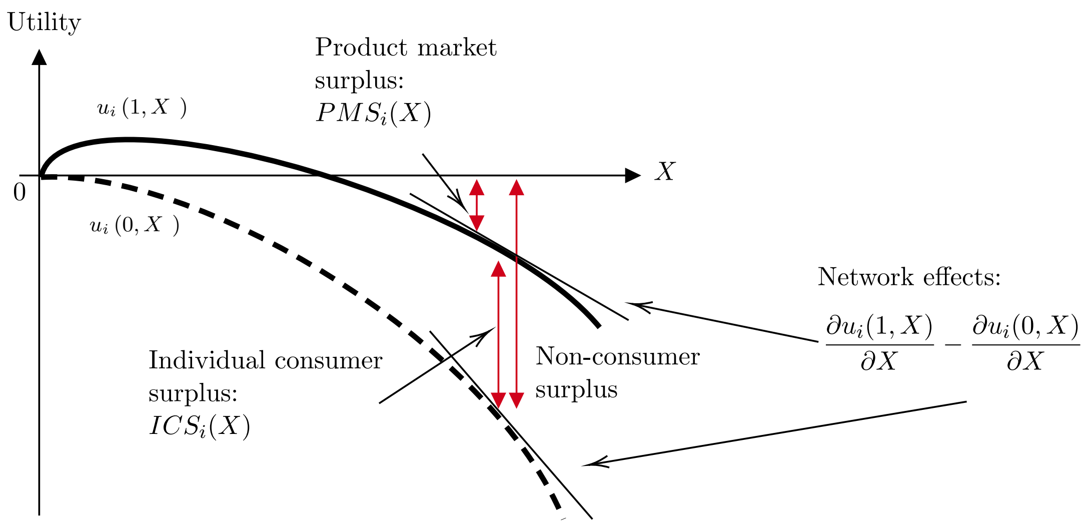

Notes: Figure A1 presents an example of individual utilities as a function of aggregate use X, as well as individual consumer surplus, product market surplus, and network effects. In this example there are negative consumption spillovers (to both users and non-users), negative product market surplus, positive individual consumer surplus, and positive network effects.

Emergence of Product Market Traps. The introspective equilibrium notion of Akerlof et al. (2023) helps shed light on how positive network effects and early adopters are essential for equilibria with product market traps to arise.[47] In such equilibrium, the quantity consumed is the limit of a sequential process where, in each stage, consumers update their beliefs about the aggregate consumption based on the aggregate consumption in the previous stage. Concretely, given some initial belief X_0, the consumption at stage $k \geq 1$,

[47] See Granovetter (1978) for an early treatment on how a population reacts dynamically to network effects.

Appendix Figure A2: Emergence of a Product Market Trap

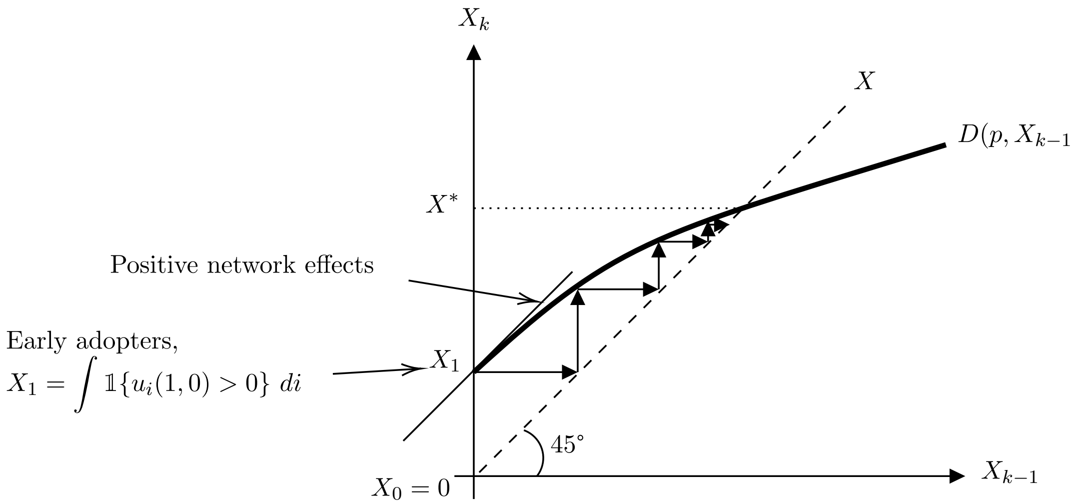

Notes: Figure A2 presents an example of the equilibrium transition given early adopters $X_1 > 0$.

X_k, is given by the aggregate demand when consumers expect a level of aggregate use equal to X_{k-1}; $X_k = D(p, X_{k-1})$. An introspective equilibrium is the limit:

$$X_1^*(p) = \lim_{k \to \infty} X_k.$$

Consider now the case of an early-stage platform without users, such that the initial expected aggregate use is zero, $X_0 = 0$. With a large enough fraction of early adopters $X_1 > 0$ that are willing to use the product when no one is using it ($u_i(1, 0) > 0$), and with positive network effects that make the demand curve upward-sloping in aggregate consumption, the expected user base in the next stage is positive, $X_1 > X_0$. The process continues, with the user base growing in every stage (due to the positive network effects) until equilibrium is reached. This transition to the equilibrium quantity X^* can happen even if a large fraction of users (including the early adopters) experience a negative utility in equilibrium ($u_i(1, X^*) < 0$). Figure A2 shows an example of this process in the case of a single equilibrium, but the same logic applies with multiple equilibria.

B Additional tables and figures

Appendix Table A1: Summary Statistics

	Obs.	Mean	Std. dev.	Median	Min	Max
Panel A: Willingness to accept (WTA) elicitations						
Panel A.1: TikTok						
Valuation Keeping Network	291	58.93	58.63	30	-10	210
Valuation Removing Network	291	39.69	52.71	30	-10	210
Product Market Valuation	291	-27.97	96.48	-10	-210	210
Product Market Valuation (with non-users)	595	-47.71	99.65	-30	-210	210
Panel A.2: Instagram						
Valuation Keeping Network	230	46.87	54.53	30	-10	210
Valuation Removing Network	230	35.57	50.57	10	-10	210
Product Market Valuation	230	-10.17	107.08	10	-210	210
Product Market Valuation (with non-users)	255	-12.98	107.32	10	-210	210
Panel A.3: Navigation/maps apps						
Valuation Keeping Network	252	51.11	53.73	30	-10	210
Valuation Removing Network	252	43.89	56.19	30	-10	210
Product Market Valuation	252	15.63	100.67	30	-210	210
Panel B: Comprehension checks						
Panel B.1: TikTok						
% Regretted elicited preferences	1,174	18.82	39.11	0	0	100
% Passed attention checks	1,174	63.97	48.03	100	0	100
Panel B.2: Instagram						
% Regretted elicited preferences	455	11.87	32.38	0	0	100
% Passed attention checks	455	63.52	48.19	100	0	100
Panel B.3: Navigation/maps apps						
% Regretted elicited preferences	493	14.81	35.55	0	0	100
% Passed attention checks	493	60.45	48.95	100	0	100
Panel C: Sample demographics						
Panel C.1: TikTok						
% Active user	595	48.91	50.03	0	0	100
% Female	569	69.42	46.12	100	0	100
Age	595	20.92	2.05	21	18	30
Panel C.2: Instagram						
% Active user	255	90.20	29.80	100	0	100
% Female	248	72.18	44.90	100	0	100
Age	255	20.86	2.14	21	18	30
Panel C.3: Navigation/maps apps						
% Active user	252	100.00	0.00	100	100	100
% Female	249	73.49	44.23	100	0	100
Age	252	20.87	2.02	21	18	30

Notes: The table presents summary statistics across all platforms, TikTok, Instagram, and navigation/maps applications. The data collection for TikTok took place in July and the data collection for Instagram and navigation/maps apps took place in August in a cross-randomized survey. The statistics depicting % of respondents are derived from dummy variables multiplied by 100. The % active user represents the fraction of respondents in the final sample who have used the platform at least once in the past month, after filtering those who do not wish to participate in the study and applying regret and attention checks. While presenting gender statistics, non-binary respondents were removed from the data in order to present the gender-binary distribution in percentage values.

Appendix Table A2: Effect of Consumption Spillovers on Welfare Estimates

	Consumer surplus			Negative surplus		
	(1)	(2)	(3)	(4)	(5)	(6)
Instagram	-4.358	-2.115		0.054*	0.049*	
	(4.927)	(3.798)		(0.030)	(0.029)	
Product Market Valuation	-35.476***	-35.476***	-35.476***	0.202***	0.202***	0.202***
	(5.902)	(5.905)	(5.890)	(0.028)	(0.028)	(0.028)
Instagram × Product Market Valuation	-21.567***	-21.567***	-21.567***	0.132***	0.132***	0.132***
	(8.318)	(8.323)	(8.301)	(0.044)	(0.044)	(0.044)
Uber Valuation		0.691***			-0.001***	
		(0.073)			(0.000)	
Dep. var. mean	26.20	26.20	26.20	0.24	0.24	0.24
Dep. var sd	86.24	86.24	86.24	0.43	0.43	0.43
Observations	964	964	964	964	964	964
Individual controls	Yes	Yes	No	Yes	Yes	No
Individual FEs	No	No	Yes	No	No	Yes

Notes: The table presents Difference-in-Differences (DiD) coefficient estimates, comparing the elicited individual and product market surplus across two platforms: Instagram and navigation/maps applications. The two dependent variables are (i) the quantitative measure of consumer surplus, denoted as *Consumer surplus*, and (ii) the direction of the consumer surplus, denoted as *Negative surplus*, represented by a binary variable coded as 1 if the surplus is negative and 0 otherwise. Columns 1 and 4 include the following individual control variables: age, gender, and the frequency of platform use, which is determined through a set of qualitative questions. Columns 2 and 5 additionally control for the valuation of the practice good, Uber. Columns 3 and 6 include individual fixed effects. Respondents who agree with their elicited valuations and those who pass all of the attention checks are included. Standard errors are given in parentheses. Standard errors are clustered at the individual level. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Appendix Table A3: Correlates of Consumer Surplus

	Valuation Keeping Network (1)	Valuation Removing Network (2)	Product Market Valuation (3)
Panel A: TikTok			
Age	2.721	2.722	2.253
	(1.840)	(1.682)	(3.410)
Female	-10.695	-11.620	-1.673
	(8.220)	(7.915)	(15.296)
Daily usage	23.578***	27.703***	13.085
	(6.866)	(5.981)	(12.359)
Dep. var. mean	58.93	39.69	-27.97
Dep. var sd	58.63	52.71	96.48
Observations	291	291	291
Panel B: Instagram			
Age	0.849	1.052	3.248
	(2.392)	(2.342)	(4.852)
Female	-0.665	-3.377	-12.966
	(8.863)	(8.609)	(19.203)
Daily usage	12.366	10.985	-0.476
	(7.930)	(7.968)	(15.819)
Network size	0.091	0.076	0.238
	(0.136)	(0.127)	(0.328)
Dep. var. mean	46.35	35.52	-6.80
Dep. var sd	54.39	50.47	106.65
Observations	181	181	181

Notes: The table presents coefficient estimates from OLS regressions. Panel A displays the results for TikTok and Panel B displays the results for Instagram. The dependent variables are Individual Consumer Surplus (*ICS*), Individual Consumer Surplus alone (*ICS alone*) and Collective Consumer Surplus (*CCS*). Columns 1-3 correspond to the elicitations in steps 1-3 in our survey, respectively. The independent variables are age, dummy variables for identifying as female and self-reported daily platform usage, and self-reported fraction of college students who are mutual friends on Instagram, labeled *Network size* in the table. *Network size* is only available in the Instagram survey and contains missing observations. Respondents who agree with their elicited valuations and those who pass all of the attention checks are included. Standard errors are given in parentheses. Standard errors are clustered at the individual level. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Appendix Figure A3: Reasons for Unwillingness to Participate in Deactivation Study

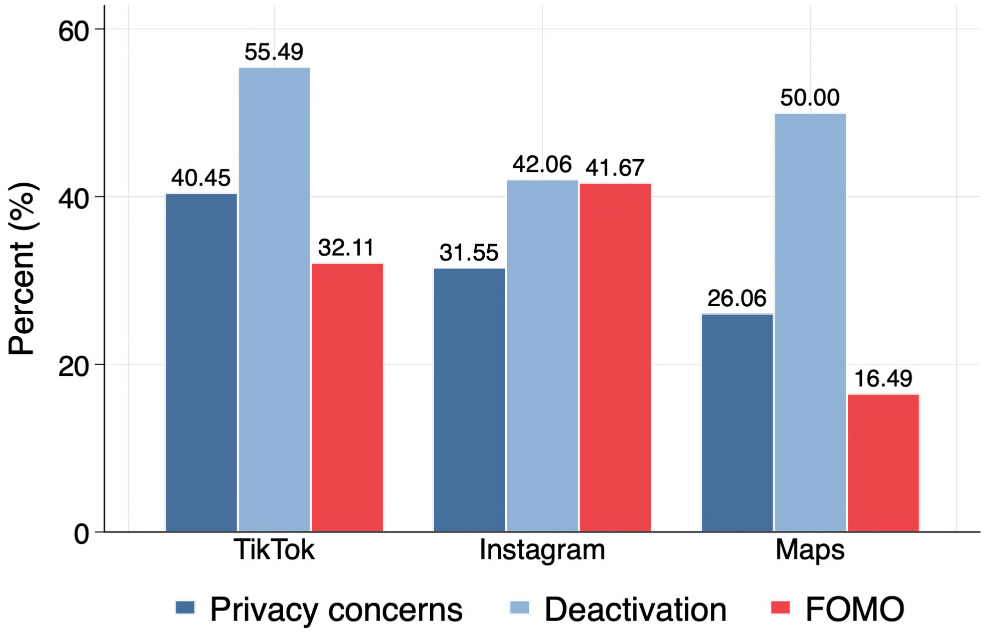

Notes: The figure presents data on the motives behind people's unwillingness to participate in the deactivation study. The respondents who declined participating in the study were asked the following question: "Why were you unwilling to participate in the study? Please select all that apply." The figure displays the fraction of respondents that were unwilling to participate because they (i) had privacy concerns, (ii) were unwilling to deactivate their account, and (iii) had a fear of missing out (FOMO). An additional "Other reason not listed above" option was available to ensure genuine feedback. As multiple selections were allowed, the categories presented above are not mutually exclusive.

Appendix Figure A4: Inverse Demand Function Across Welfare Measures

(a) TikTok

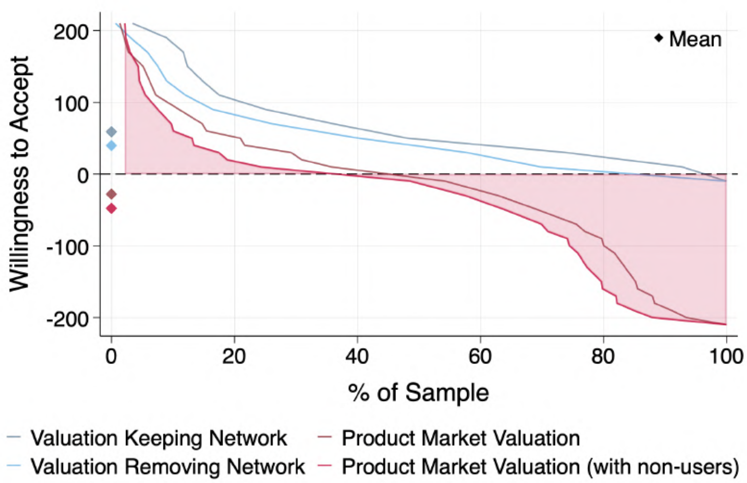

(b) Instagram

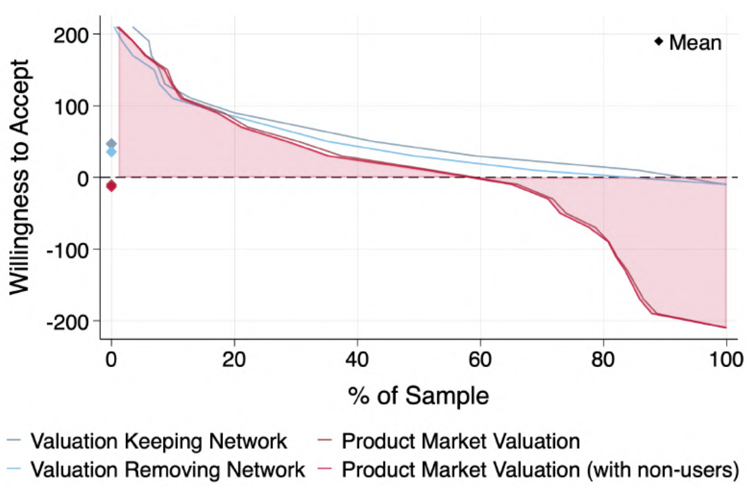

Notes: Figure A4 displays the inverse demand function of respondents' valuation for our different welfare measures. Panel (a) presents the results for TikTok and Panel (b) presents the results for Instagram. Respondents who agree with their elicited valuations and those who pass all of the attention checks are included.

Appendix Figure A5: Distribution of Consumer Surplus Across Welfare Measures

(a) TikTok

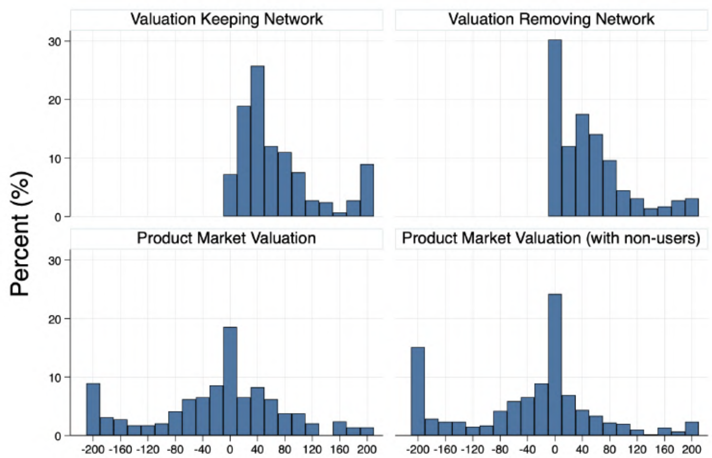

(b) Instagram

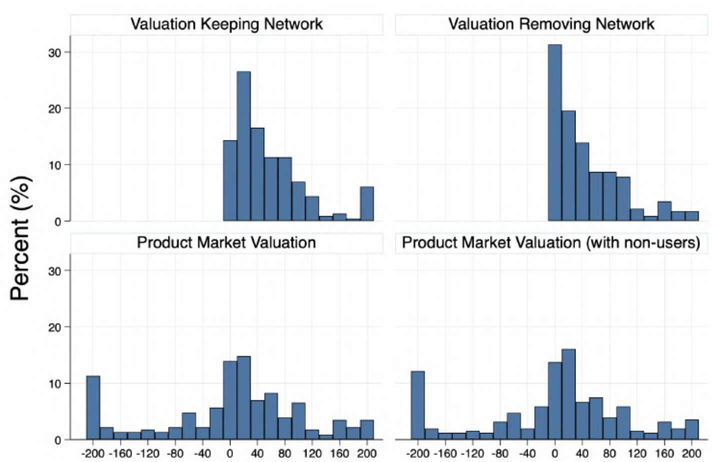

Notes: Figure A5 presents the probability density function of valuations for the different welfare measures. Panel (a) presents the results for TikTok and Panel (b) present the results for Instagram. Respondents who agree with their elicited valuations and those who pass all of the attention checks are included.

Appendix Figure A6: Hypothetical Ranking of Alternatives

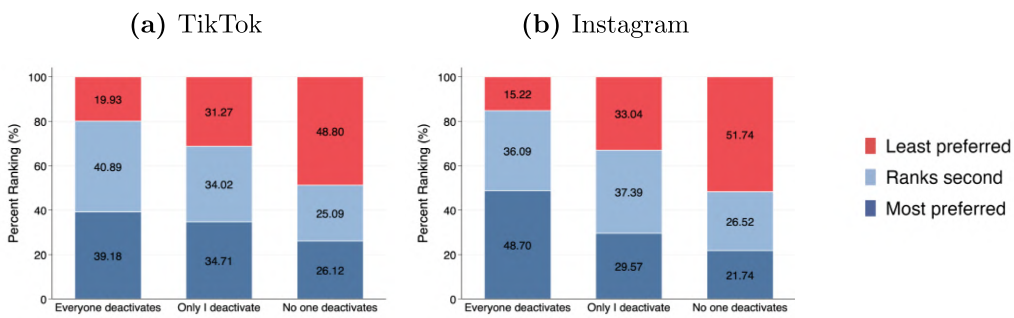

Notes: Figure A6 presents participants' ranking of three hypothetical scenarios about the deactivation of the social media platform: (i) Everyone deactivates (ii) Only I deactivate (iii) No one deactivates. Panel (a) displays results for TikTok and Panel (b) shows results for Instagram. The area in dark blue indicates people's most preferred option; the area in light blue indicates the option ranked second; the area in red shows the least preferred option.

Appendix Figure A7: Validation of Hypothetical Survey Questions

TikTok **Instagram**

(a) World With/Without

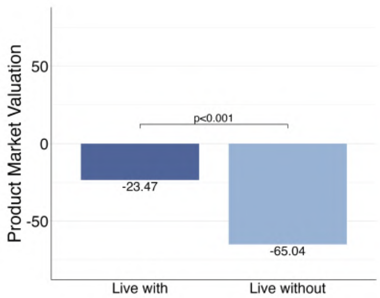

(b) World With/Without

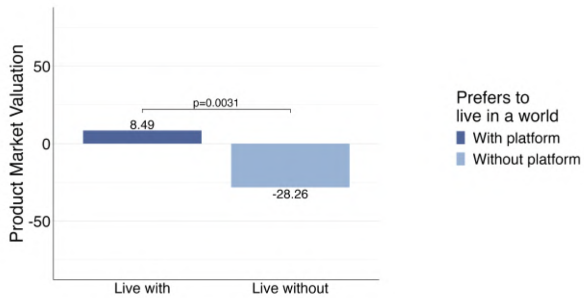

(c) Hypothetical Ranking

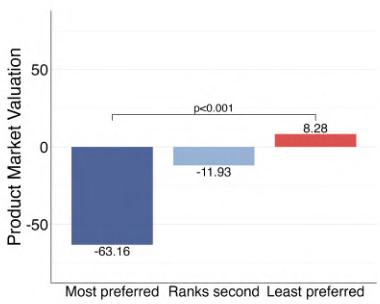

(d) Hypothetical Ranking

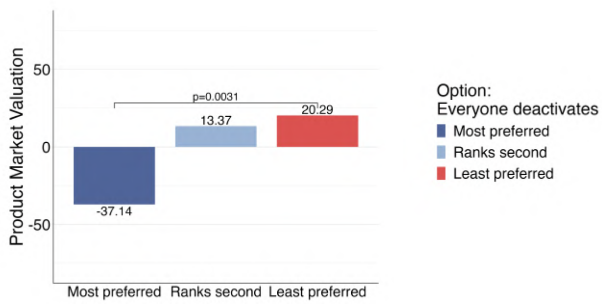

(e) Perceived Likelihood of Large-Scale Deactivation Taking Place

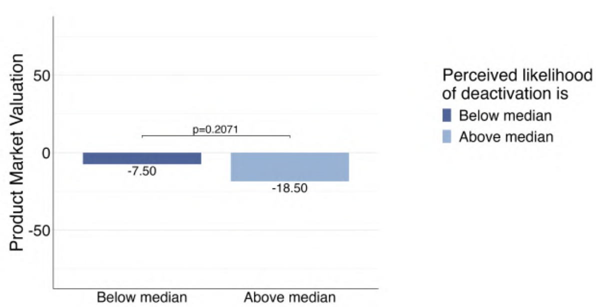

Notes: Figure A7 presents a validation of the hypothetical survey questions. The outcome variable in the figures is the *Product Market Valuation*. Panels (a) and (c) show results for TikTok. Panels (b), (d) and (e) present results for Instagram. Panels (a) and (b) present the *Product Market Valuation* by people's preference to live in a world with or without the platform. Panels (c) and (d) presents the *Product Market Valuation* by people's hypothetical ranking of the deactivation for everyone. Panel (e) presents the *Product Market Valuation* by respondents' perceived likelihood of the large-scale deactivation study taking place.

Appendix Figure A8: Regretters Across Steps

(a) TikTok

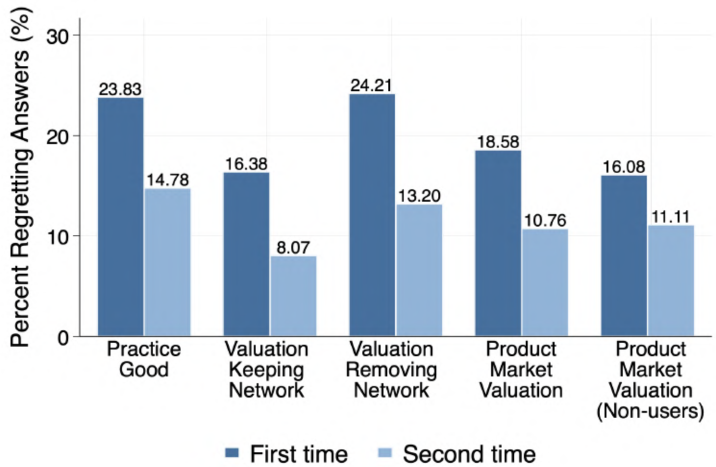

(b) Instagram

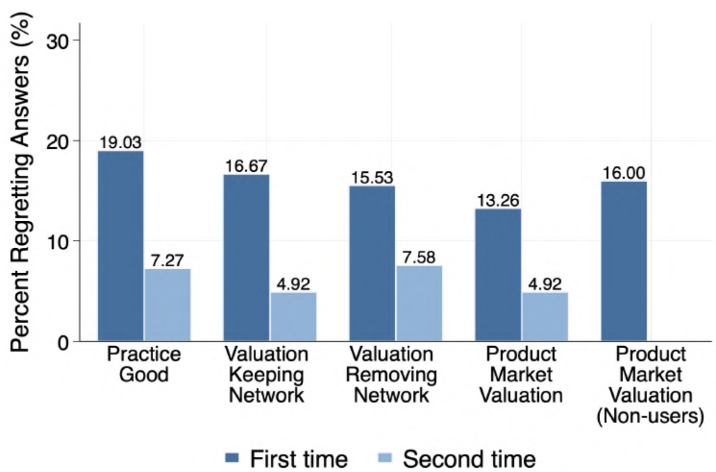

Notes: Figure A8 presents the fraction of respondents that regret their choices across the different measures. Panel (a) presents the results for TikTok and Panel (b) present the results for Instagram. Dark blue bars indicate the fraction of respondents regretting their choices the first time they completed a given valuation. Light blue bars indicate the fraction of respondents regretting their choices the second time they completed a given valuation.

Appendix Figure A9: Consumer Welfare: Navigation and Maps Smartphone Apps

(a) Consumer Surplus Across Welfare Measures

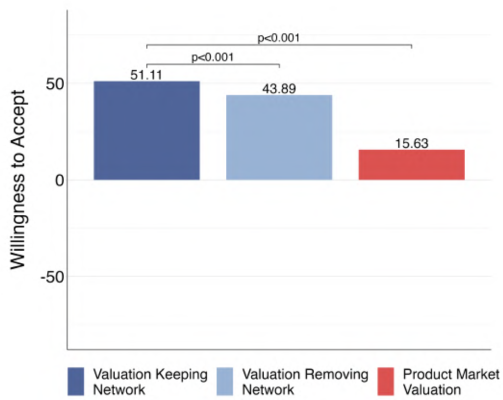

(b) Fraction of Respondents Deriving Negative Welfare

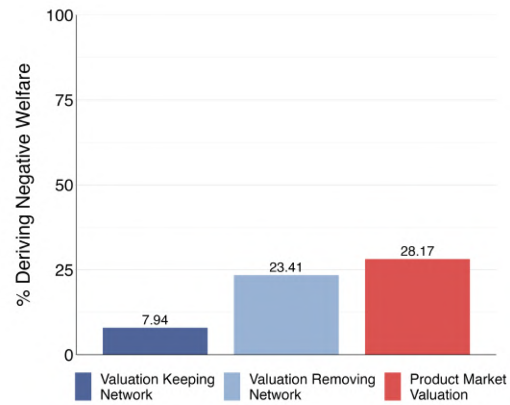

(c) Distribution of Consumer Surplus Across Welfare Measures

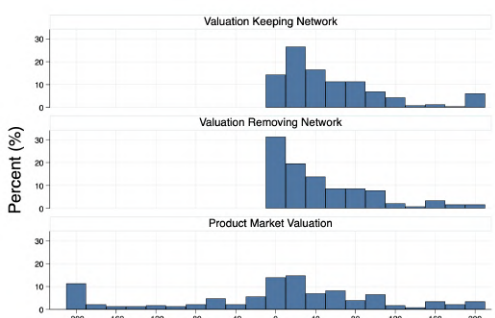

(d) Ranking

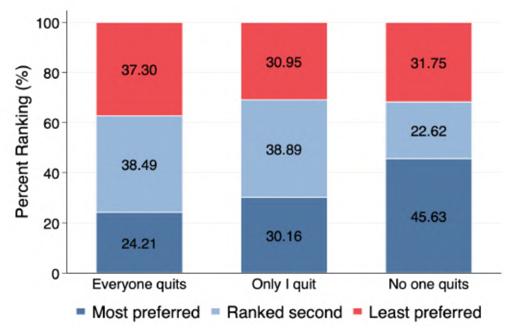

Notes: Figure A9 presents the survey results for navigation and maps smartphone apps. Figure A9a average valuations for the different welfare measures. Figure A9b presents the fraction of users with negative welfare across the different welfare measures for navigation and maps smartphone apps. Figure A9c presents the probability density function of valuations for the different welfare measures for navigation and maps smartphone apps. Figure A9d presents participants' responses ranking of three hypothetical scenarios: (i) All participating students quit using navigation apps (ii) Only I quit using navigation apps (iii) No one quits using navigation apps. The area in dark blue indicates people's most preferred option; the area in light blue indicates the option ranked second; the area in red shows the least preferred option. In all figures, respondents who agree with their elicited valuations and those who pass all of the attention checks are included.

B.1 Additional Robustness Checks

B.1.1 Sample Selection

Appendix Figure A10: Consumer Surplus Across Welfare Measures: Full Sample

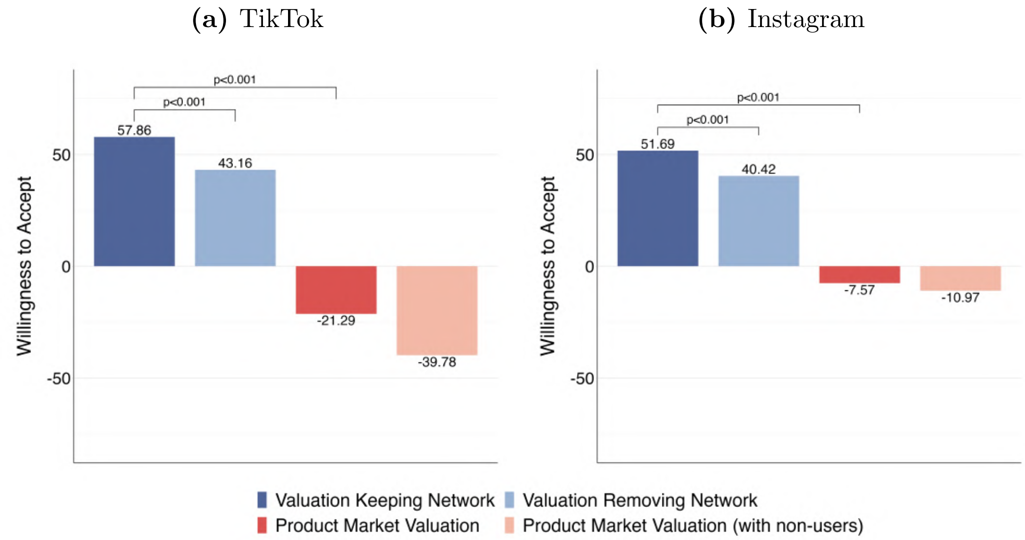

Notes: Figure A10 presents average valuations for the different welfare measures. Panel (a) presents the results for TikTok and Panel (b) present the results for Instagram. Inattentive respondents and respondents regretting their choices are also included in this specification. The first three bars in each panel represent valuations exclusively for active users. The fourth bar represents the average valuation of active users and non-users. Reported p-values correspond to one-sided t-tests testing the null hypothesis that individual welfare estimates are lower than the aggregate welfare estimate.

Appendix Figure A11: Consumer Surplus Across Welfare Measures: Excluding Regretters and Including Inattentive Respondents

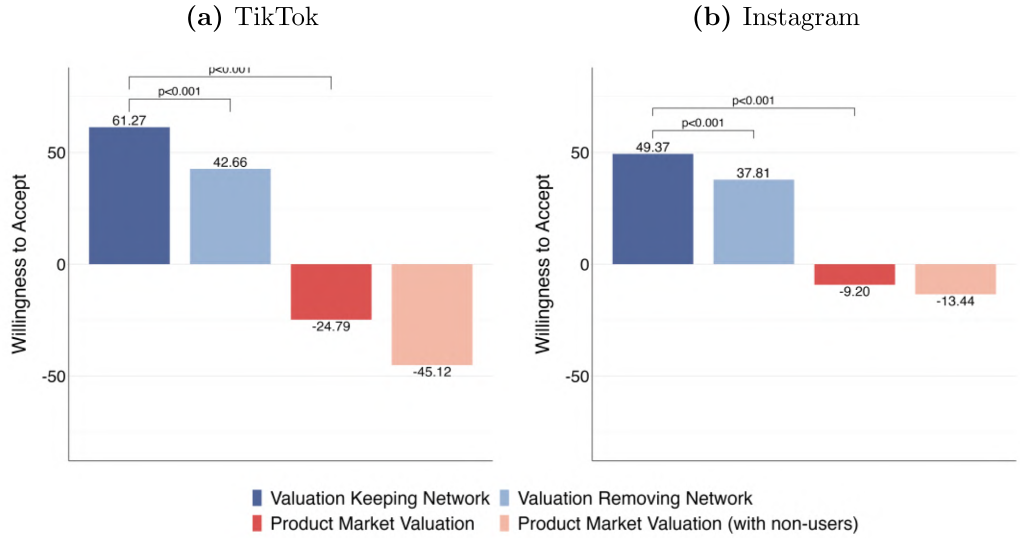

Notes: Figure A11 presents average valuations for the different welfare measures. Panel (a) presents the results for TikTok and Panel (b) present the results for Instagram. Only respondents who regretted any of their choices are excluded, while inattentive respondents are included. The first three bars in each panel represent valuations exclusively for active users. The fourth bar represents the average valuation of active users and non-users. Reported p-values correspond to one-sided t-tests testing the null hypothesis that individual welfare estimates are lower than the aggregate welfare estimate.

Appendix Figure A12: Consumer Surplus Across Welfare Measures: Excluding Inattentive Respondents and Including Regretters

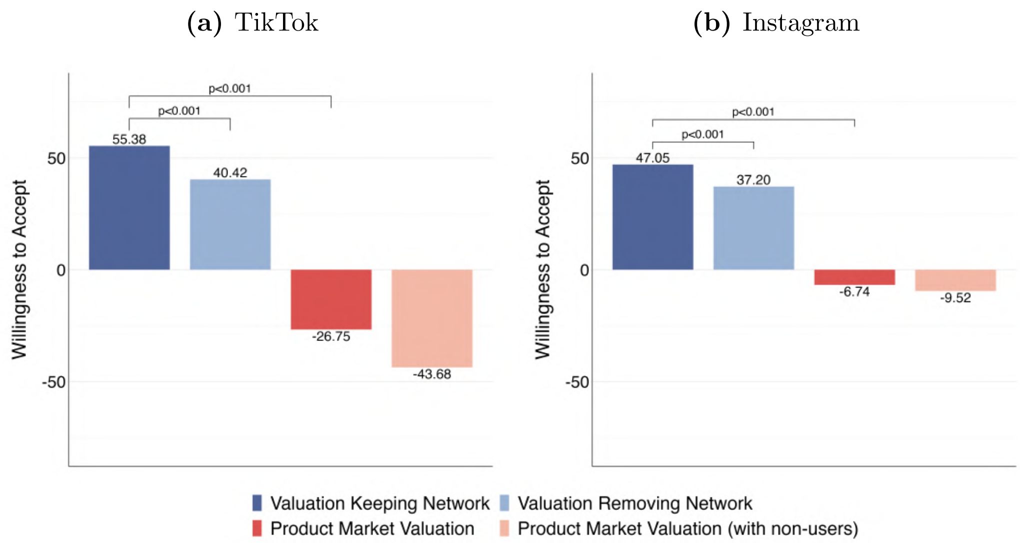

Notes: Figure A12 presents average valuations for the different welfare measures. Panel (a) presents the results for TikTok and Panel (b) present the results for Instagram. Only respondents who failed to pass all attention checks are excluded, while those who regretted their choices are included. The first three bars in each panel represent valuations exclusively for active users. The fourth bar represents the average valuation of active users and non-users. Reported p-values correspond to one-sided t-tests testing the null hypothesis that individual welfare estimates are lower than the aggregate welfare estimate.

B.1.2 Triangular Distribution Estimates

In our primary analysis, we employ a multiple price list to narrow down the WTA range of our respondents to $20 and assign the mean of each respondent's lower and upper bounds to obtain a unique WTA. As a robustness check, we employ an alternative distributional assumption: a triangular distribution to account for potential biases stemming from the unbounded nature of our lowest and highest intervals.

Following Allcott and Kessler (2019) we assume a triangular distribution at the unbounded ranges. To determine a new upper bound, we compute the mass at the upper unbounded interval, [$200, $220], and the density at the preceding interval, [$180, $200]. Then, using the formula for the probability density function (PDF) for a triangular distribution, we determine the alternative upper bound of the distribution. Subsequently, we compute the mean for the upper unbounded range. Analogously, for the lower unbounded interval, using the same principles we determine a new lower bound and substitute it with $-10. Figure A13 shows the willingness to accept means for each category, estimated assuming a triangular distribution.

Appendix Figure A13: Consumer Surplus Across Welfare Measures: Triangular Distribution Estimates

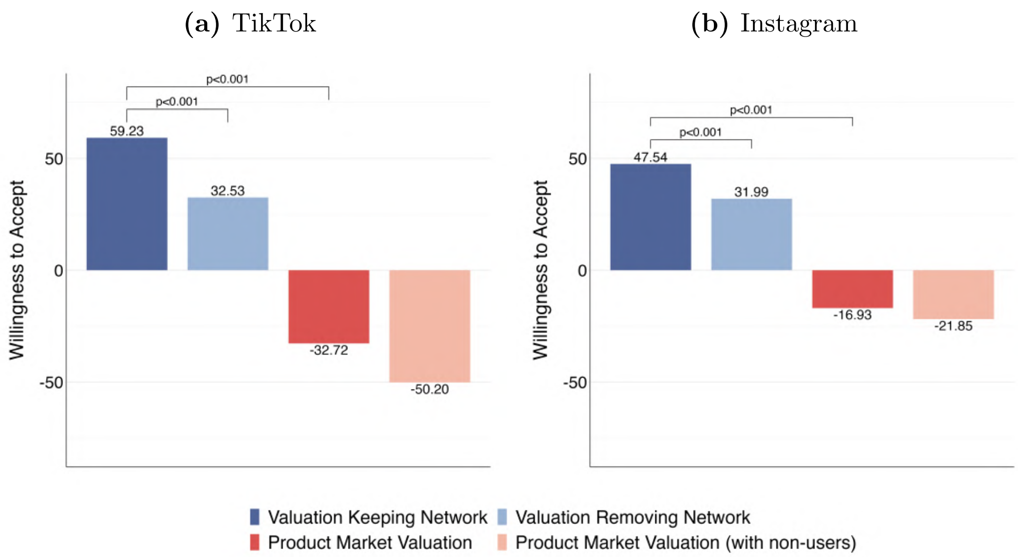

Notes: Figure A13 presents average valuations for the different welfare measures assuming triangular distributions for unbounded intervals. Panel (a) presents the results for TikTok and Panel (b) present the results for Instagram. Respondents who agree with their elicited valuations and those who pass all of the attention checks are included. The first three bars in each panel represent valuations exclusively for active users. The fourth bar represents the average valuation of active users and non-users. Reported p-values correspond to one-sided t-tests testing the null hypothesis that individual welfare estimates are lower than the aggregate welfare estimate.

B.2 Open ended responses

Our surveys included two open-ended questions to provide direct evidence on the mechanisms and motives driving consumption. The purpose of this section is two-fold; first, we present an overview of the hand-coding schemes we employed for the categorization of open-ended responses. Second, we summarize the validation of our manual hand-coding with artificial intelligence methods, presenting results from both techniques.

Hand-coding schemes. Table A4 presents the hand-coding scheme applied to open-ended responses for the question: "You mentioned you would prefer to live in a world without [platform]. Why do you still use it?". As implied by the phrasing, this question targeted only those respondents who previously expressed a desire to live without the platform. Table A4 presents the hand-coding scheme used for the question: "How would you feel if you were the only one who quit using [platform] and everyone else kept using it?". Note that the categories "Negative", "Beneficial", and "Other", encompass several subcategories. Specifically, "Negative" includes responses mentioning unfairness, impracticality, feeling inferior, dependent, bad, stressed, or lost; whereas "Beneficial" includes responses mentioning self-improvement and feeling positively challenged, unpressured, or good.

Certain respondents expressed a conditional indifference based on compensation or the deactivation's duration. Given that this does not truly signify 'indifference', such responses were categorized under "Other".[48] The category also includes relatively infrequent subcategories such as a stated preference to deactivate themselves in order to prevent inconvenience to others, and deriving satisfaction from going against the norm. For both open-ended questions, responses that were non-sensical were excluded from the analysis (N=11, 1.88%).

Validation with Artificial Intelligence. To corroborate our manual categorization, we employed recent artificial intelligence methods, in particular a powerful large language model (GPT-4). We structured a validation exercise with the prompt: "You will be supplied with a list of responses. The responses refer to the usage of different platforms, the platform will be indicated in parentheses at the end of the response. Please classify responses based on the coding scheme below. Please note that each open-ended response can fall into multiple categories or even none." Subsequent to this, we supplied GPT-4 with the hand-coding scheme, complete with category names, definitions, and illustrative examples. To maintain methodological consistency between our manual coding and GPT-4's process, we provided GPT-4 with definitions and examples for each subcategory in the subsequent question. These subcategories were subsequently grouped under the primary categories.

Figure A14 displays the category distributions by platform and coding methods. Panel A presents the results for the open-ended question aimed at eliciting the motives for social media consumption despite a preference to live in a world without it; while Panel B

[48] These open-ended questions followed the willingness-to-accept elicitation questions, potentially leading some respondents (10%, N=59) to mistakenly believe that the "deactivation" pertained to the study's deactivation, which was in exchange for monetary payment and had a four-week duration.

presents the results for second question aimed at unraveling the mechanisms behind non-user consumption spillovers. The juxtaposition of the results of the two coding methods demonstrates that both methods yield remarkably similar results.

To further validate our hand-coding, we conduct a correlational exercise for each category. Once again, the results are presented per question. Each column represents the categories employed for the coding schemes. As displayed in Table A6, all categories have large and statistically significant correlation coefficients across the two methods.

Appendix Table A4: Overview of hand-coding scheme for reasons to use TikTok/Instagram/Maps despite preferring a world without it

Category	**Definition**	**Example(s)**
FOMO	Respondent mentions fear of missing out, feeling out of the loop, their wish to stay connected, or justifies usage through others' usage	"I feel compelled to keep 'in touch' with what I perceive as being the culturally relevant 'thing' at the moment. It breeds a sense of FOMO when you don't use it." (*TikTok*); "Everyone else uses it so I feel that I will be missing out if I don't." (*Instagram*); "I still use navigation maps because it is what everyone uses [...]" (*Maps*)
Entertainment	Respondent mentions they use it to be entertained	"It's a very good source of entertainment and it's always something to do when bored." (*TikTok*); "It's a default way to pass time when I'm bored." (*Instagram*);
Addiction	Respondent mentions inability to let go or directly mentions addiction	"I use TikTok as a habit. I hate TikTok and know that I have other things I need to do, but I subconsciously click on it, then scroll for hours. It's very hard to control it." (*TikTok*); "Because I am addicted to the scrolling and tired of wasting valuable time on the app." (*Instagram*)
Information	Respondent mentions informational purposes such as following the news, keeping abreast of college events, or getting directions.	"for information on current events because i do not watch the news" (*TikTok*); "I use it to keep inform about my university events and news" (*Instagram*); "I don't know where to go" (*Maps*)
Productivity/ Convenience	Respondent mentions convenience of use or states to use platform for productive/business purposes.	"It's easy to see stuff I like (art, new art news, movie reviews, etc)." (*TikTok*); "I still use instagram for business purposes." (*Instagram*); "It's more convenient than pulling out a map and I have a terrible sense of direction" (*Maps*)

Notes: The table displays an overview of the hand-coding scheme used for categorizing the open-ended answers given to the question: "You mentioned you would prefer to live in a world without TikTok/Instagram/navigation apps. Why do you still use it/them?". The question was only asked to participants that are are active users of the respective platforms and stated they would prefer to live in a world without said platform. Respondents who agree with their elicited valuations and those who pass all of the attention checks are included.

Appendix Table A5: Overview of hand-coding scheme for how respondent would feel if only they were to quit using platform

Category	**Definition**	**Example**
FOMO	Respondent mentions fear of feeling missing out, left out, or being out of the loop	"I would probably feel somewhat out of the loop when it comes to trends, with a consistent feeling of FOMO." (*TikTok*); "I would feel really left out since a lot of people use it to communicate about events and parties and with one another" (*Instagram*); "I would feel a bit isolated, maybe excluded from certain conversations involving travel plans, etc" (*Maps*)
Negative	Respondent expresses negative emotions; that it would be unfair, impractical, etc.	"it would be a little unfair" (*TikTok*); "[...] feel discouraged and jealous of everyone else." (*Instagram*); "I would feel lost and not confident in my ability to navigate" (*Maps*)
Indifferent	Respondent states that they would not be particularly affected	"No different, because I don't use tiktok often anyway"; (*TikTok*); "I would be fine. I don't really post on Instagram. It wouldn't be much of a change." (*Instagram*); "Wouldn't mind as long as knew my way around" (*Maps*)
Beneficial	Respondent mentions deriving a benefit	"relieved, probably." (*TikTok*); "I would feel free"; (*Instagram*); "It will be an awesome experiment and experience, asking everyone for directions" (*Maps*)
Other	Diverse set of motives; including substituting platform, indifference conditional on getting paid, or fondness to spare others the struggle	"I would just use other social media" (*TikTok*); "It's okay as long as I have some monetary benefit in it."; (*Instagram*); "... I don't want everyone else to struggle especially since people have different like circumstances" (*Maps*)

Notes: The table displays an overview of the hand-coding scheme used for categorizing the open-ended answers given to the question: "How would you feel if you were the only one who deactivated (quit using) TikTok/Instagram (navigation/maps apps) and everyone else kept using it (them)?". The question was only asked to participants that are active users of the respective platform. In the TikTok survey, the question was further restricted to respondents who stated they would prefer to live in a world without TikTok. We did not apply this restriction for the Instagram/maps survey. Respondents who agree with their elicited valuations and those who pass all of the attention checks are included.

Appendix Figure A14: Validation based on Large Language Model

(a) Motives for social media consumption despite a preference to live in a world without it

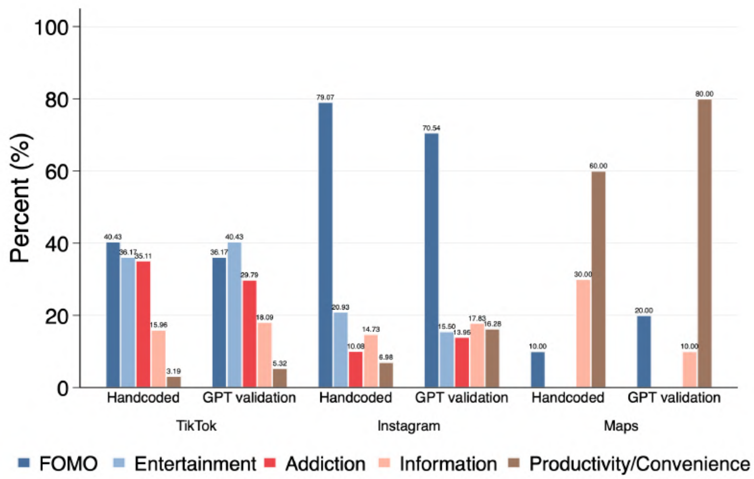

(b) Mechanisms behind non-user consumption spillovers

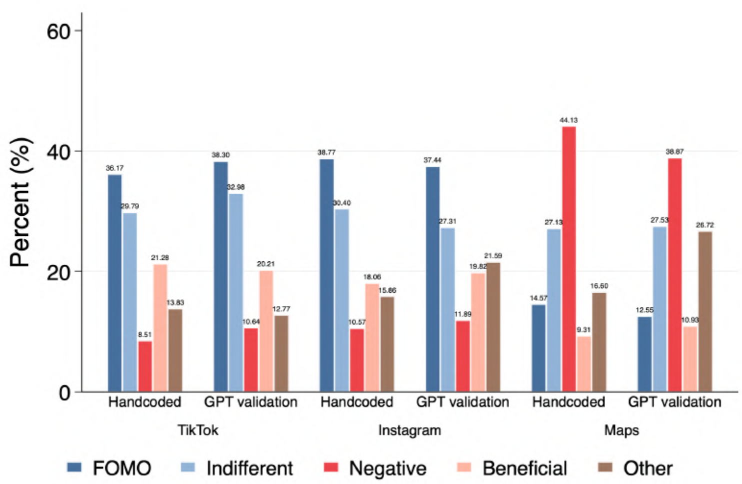

Notes: Figure A14 presents the distribution of categories based on open-ended responses separately for the hand-coded data and the data coded by a large language model (GPT4). Panel (a) details the results for the question, "You mentioned you would prefer to live in a world without TikTok/Instagram/navigation apps. Why do you still use them?" Meanwhile, Panel (b) showcases the results for the question, "How would you feel if you were the only one who stopped using TikTok/Instagram/navigation apps, while everyone else continued their use?

Appendix Table A6: Validation of hand-coded data from Large Language Model

Panel A: Motives for social media consumption despite a preference to live in a world without it					
	FOMO	Entertainment	Addiction	Information	Productivity/ Convenience
Correlation coefficient	0.814	0.718	0.837	0.724	0.563
	(0.038)	(0.046)	(0.036)	(0.045)	(0.054)
Hand-coded responses:					
Mean	0.605	0.262	0.197	0.159	0.077
Std. dev.	0.490	0.441	0.399	0.366	0.268
GPT-4 coded responses:					
Mean	0.545	0.249	0.197	0.176	0.146
Std. dev.	0.499	0.433	0.399	0.382	0.354
Observations	233	233	233	233	233

Panel B: Evidence on mechanisms behind non-user consumption spillovers					
	FOMO	Indifferent	Negative	Beneficial	Other
Correlation coefficient	0.894	0.815	0.674	0.697	0.658
	(0.019)	(0.024)	(0.031)	(0.030)	(0.032)
Hand-coded responses:					
Mean	0.278	0.289	0.248	0.148	0.158
Std. dev.	0.448	0.454	0.432	0.355	0.365
GPT-4 coded responses:					
Mean	0.268	0.283	0.234	0.160	0.224
Std. dev.	0.443	0.451	0.424	0.367	0.417
Observations	568	568	568	568	568

Notes: The table presents the correlation coefficients between our manual categorization and the GPT-4 categorization of open-ended responses. Each column corresponds to a specific category used in the classification process. Correlation coefficients were calculated using dummy variables: for every coding technique, a dummy variable is set to 1 if the open-ended response fits within a particular category. These coefficients then show the correlation between these dummy variables. Panel A details the results for the question, "You mentioned you would prefer to live in a world without TikTok/Instagram/navigation apps. Why do you still use them?" Meanwhile, Panel B showcases the results for the question, "How would you feel if you were the only one who stopped using TikTok/Instagram/navigation apps, while everyone else continued their use?" Standard errors are given in parentheses and are computed based on the Pearson correlation coefficient formula.

C Experimental Instructions

C.1 TikTok: July 2023

C.1.1 Introduction to Survey and Deactivation Study Instructions

University of Chicago
Online Consent Form for Research Participation
Study Number: IRB23 - 0797
Researcher: Leonardo Bursztyn

Description: You will be asked to fill out a short survey. Participation is voluntary and takes about 8 minutes.

Incentives: Upon completion of this survey, you will be compensated by your survey provider.

Risks and Benefits: There are no foreseeable risks associated with this study beyond those involved in answering a survey. The research team cannot and do not guarantee or promise that you will receive any benefits from this study.

Confidentiality: Confidentiality of your research records will be strictly maintained by storing any personally identifiable data in secure accounts that can only be accessed by researchers in this study. After the experiment is over, we will delete any personally identifiable information from our dataset and replace it with an arbitrary participant number. This will allow us to maintain your privacy in all published and written data resulting from the study. Information not containing identifiers may be used in future research or shared with other researchers without your additional consent. If you decide to withdraw, data collected up until the point of withdrawal may still be included in the analysis.

Contacts & Questions: If you have questions or concerns about the study, you can contact the researcher at: bursztyn@uchicago.edu.
If you have any questions about your rights as a participant in this research, feel you have been harmed, or wish to discuss other study-related concerns with someone who is not part of the research team, you can contact the University of Chicago Social & Behavioral Sciences Institutional Review Board (IRB) Office by phone at (773) 702-2915, or by email at sbs-irb@uchicago.edu

Consent: Participation is voluntary. Refusal to participate or withdrawing from the research will involve no penalty or loss of benefits to which you might otherwise be entitled.
By clicking "Agree" below, you confirm that you have read the consent form, are at least 18 years old, and agree to participate in the research. Please print or save a copy of this page for your records.

○ I **agree** to participate in the research. I confirm that I am above 18 years of age or older.

○ I **do NOT agree** to participate in the research. You will be directed to an exit screen.

This survey is directed at university students' social media preferences. It is designed by our research team at the University of Chicago and is administered via partnership with College Pulse.*

Thank you for participating!

*College Pulse is a research and analytics company that specifically aims to understand the attitudes, preferences, and behaviors of today's college students.

How frequently did you use each of the following **social media platforms** in the past month?

	Not at all	Once	Once a week	Twice a week	Every day
Instagram	○	○	○	○	○
Facebook	○	○	○	○	○
TikTok	○	○	○	○	○
Twitter	○	○	○	○	○

Please read the instructions of this survey carefully. We will ask you some questions to ensure your understanding of its content and give you a **bonus payment** if you answer correctly!

TikTok

As mentioned, this survey is directed at university students' social media preferences.

After this survey, we will conduct a study in which we will ask **students at your university to deactivate their TikTok accounts** for **four weeks** in exchange for a **monetary payment.**

"Deactivation" studies like this have been conducted in the past (e.g., by Mosquera et al., 2018 and Allcott et al., 2020) with close to 90% compliance.

Should students deactivate their TikTok, they can go back to using it whenever they want, with their content and network unchanged, but they will then forgo any monetary payment.

To verify that students deactivate their TikTok accounts, we will **periodically visit their profiles** and **require them to upload screenshots of their app usage.**

The next part of the study involves asking you if you want to deactivate your TikTok account for four weeks in exchange for monetary payment.

For that we will need to collect your TikTok handle and we would ask you to submit periodic screenshots of your phone's time use statistics if you are selected.

Would you be willing to participate?

- ○ Yes
- ○ No

TikTok

How will we verify that selected users deactivate their TikTok accounts?

- ○ By asking them to install an app
- ○ By periodically visiting their profiles
- ○ By requiring them to upload screenshots
- ○ By periodically visiting their profiles and requiring them to upload screenshots

TikTok

For how long will we ask selected users to deactivate their TikTok accounts?

- ○ Four weeks
- ○ Eight weeks
- ○ Ten weeks
- ○ One week

C.1.2 Step 0: Practice Good

Before we begin, we will ask you a series of hypothetical practice questions for you to get accustomed to our survey.

Uber

To understand how much people value their Uber account, **we ask university students, including you,** to decide whether to **deactivate their Uber account** for four weeks, in exchange for different monetary payments.

A computer will **randomly select one student from your university** to be eligible for the deactivation study.

Uber

The next questions involve real money. The computer will randomly generate an amount of money to offer you to participate in the deactivation study.

We will now ask you a series of questions offering you different payment scenarios in case you are selected for the deactivation study.

- If you accept any price scenario lower than the computer's offer, we will invite you to the deactivation study and give you the computer's offer.
- If you do not accept any price scenario lower than the computer's offer, we will not invite you to the deactivation study even if you are selected.

This rule means that the higher the amount you require the lower the chance that you will receive the computer's offer.

Therefore, while answering the following questions, **please choose carefully as each answer could be "the one that counts."**

We will now ask you a comprehension question based on the text above:

Which of the following statements is true?

- ○ The amount I require does not affect the chance that I receive the computer's offer
- ○ The higher the amount I require the higher the chance that I receive the computer's offer
- ○ The higher the amount I require the lower the chance that I receive the computer's offer

Uber

Which of the following options would you prefer?

- ○ I deactivate my Uber account AND I receive $140
- ○ I keep my Uber account active

C.1.3 Step 1: Valuation Keeping Network

You have now completed the practice section. We will now ask you a series of questions about your **social media preferences**.

To establish appropriate payment amounts for the deactivation study, **we ask university students, including you**, to decide whether to deactivate their TikTok accounts for different monetary amounts.

A computer will **randomly select one student from your university** to be eligible for the deactivation study.

The next questions involve real money. The computer will randomly generate an amount of money to offer you to participate in the deactivation study.

We will now ask you a series of questions offering you different payment scenarios in case you are selected for the deactivation study.

Therefore, while answering the following questions, **please choose carefully as each answer could be "the one that counts."**

Which of the following options would you prefer?

○ I deactivate my TikTok account AND I receive $120

○ I keep my TikTok account active

Which of the following options would you prefer?

○ I deactivate my TikTok account AND I receive $60

○ I keep my TikTok account active

Which of the following options would you prefer?

- ◯ I deactivate my TikTok account AND I receive $20
- ◯ I keep my TikTok account active

Which of the following options would you prefer?

- ◯ I deactivate my TikTok account
- ◯ I keep my TikTok account active

According to your answers to the previous questions, you would require a payment worth between $0 and $20 to deactivate your TikTok account for four weeks.

Do you agree with the above statement about your **valuation**?

- ◯ Yes
- ◯ No

C.1.4 Step 1: Valuation Removing Network

TikTok

College Pulse has a panel exceeding 650,000 university students.* We are targeting universities with a high penetration of College Pulse.

We will now ask you to consider two additional options for a large-scale deactivation of TikTok at your university. **One of them will be randomly implemented** if we manage to recruit **more than two-thirds of the students at your university.**

We expect 90% of students to comply with deactivation based on previous studies (e.g., by Mosquera et al., 2018 and Allcott et al., 2020).

*See https://collegepulse.com.

TikTok

Option 1 for the large-scale deactivation study:

In collaboration with College Pulse, we will ask students at your university sequentially whether they would like to deactivate their TikTok accounts.

Consider the scenario where we have asked **all participating students at your university** to deactivate their TikTok accounts for four weeks in exchange for a payment.

TikTok

The next questions involve real money. The computer will randomly generate an amount of money to offer you to participate in the deactivation study.

We will now ask you a series of questions offering you different payment scenarios in case you are selected for the deactivation study.

Therefore, while answering the following questions, **please choose carefully as each answer could be "the one that counts."**

Consider the scenario where we have asked **all participating students at your university** to **deactivate their TikTok accounts** for four weeks in exchange for monetary payment.

If you had the choice to **also** deactivate your TikTok account for the next four weeks, which of the following options would you prefer?

○ I deactivate my TikTok account if the other students deactivate their TikTok accounts

○ I keep my TikTok account active if the other students deactivate their TikTok accounts

According to your answers to the previous questions, you would deactivate your TikTok account for four weeks without monetary payment, if we ask all participating students at your university to deactivate their TikTok accounts.

Do you agree with the above statement about your **valuation**?

○ Yes

○ No

C.1.5 Step 3: Product Market Valuation

TikTok

Option 2 for the large-scale deactivation study:

We know how much we would need to pay **every participating student at your university** to deactivate their TikTok accounts for four weeks.

To give everyone an equal chance to decide, we will randomly choose one of the students.

This student's identity will remain anonymous, and they can choose one from the following options:

1. **We ask all participating students with a TikTok account to deactivate it**, or
2. **We keep things as they are**.

If you decide for all participating students to deactivate their TikTok accounts:

- **We will pay the other students the amount they required** and we will **establish your payment, if any, below.**
- The deactivation study will be stopped for everyone only if you go back to using TikTok before the end of the four weeks.
- If the study is stopped early, you will not receive payment and we will pay the other students based on the actual time they spend in the study.
- If someone from the other participating students goes back to using TikTok before the end of the study, they themself will not receive any payment.

Suppose that you decide for us to ask all participating students to deactivate their TikTok accounts:

Which of the following statements is correct?

- ○ We will pay the other students more than what they required to deactivate their TikTok accounts
- ○ We will force the other students to deactivate their TikTok accounts
- ○ We will pay the other students what they required to deactivate their TikTok accounts

The next questions involve real money. The computer will randomly generate an amount of money to offer you to participate in the deactivation study.

We will now ask you a series of questions with different payment scenarios in case you are selected.

Therefore, while answering the following questions, **please choose carefully as each answer could be "the one that counts."**

Which of the following options would you prefer?

- ○ All participating students at my university deactivate their TikTok accounts
- ○ All participating students at my university keep their TikTok accounts active

Which of the following options would you prefer?

- ○ All participating students at my university deactivate their TikTok accounts
- ○ All participating students at my university keep their TikTok accounts active AND I receive $140

Which of the following options would you prefer?

- All participating students at my university deactivate their TikTok accounts
- All participating students at my university keep their TikTok accounts active AND I receive $200

According to your answers to the previous questions, you would forgo a payment worth between $180 and $200 to have all participating students at your university, including you, deactivate their TikTok accounts for four weeks.

Do you agree with the above statement about your **valuation**?

- Yes
- No

C.1.6 Qualitative Questions

Please rank the following options from your most preferred (1) to your least preferred (3).

1. Every student at my university, including me, deactivates TikTok
2. No one deactivates TikTok
3. I deactivate TikTok and every other student at my university keeps using it

Would you prefer to live in a world with or without TikTok?

- I would prefer to live in a world with TikTok
- I would prefer to live in a world without TikTok

You mentioned you would prefer to live in a world without TikTok.

Why do you still use it?

How would you feel if you were the only one deactivating your TikTok and everyone else kept using it?

Now some **demographic** questions.

Which of the following describes you more accurately?

- Male
- Female
- Other / Prefer not to say

What is your age?

C.2 Instagram and Maps: August and September 2023

C.2.1 Introduction to Survey and Deactivation Study Instructions

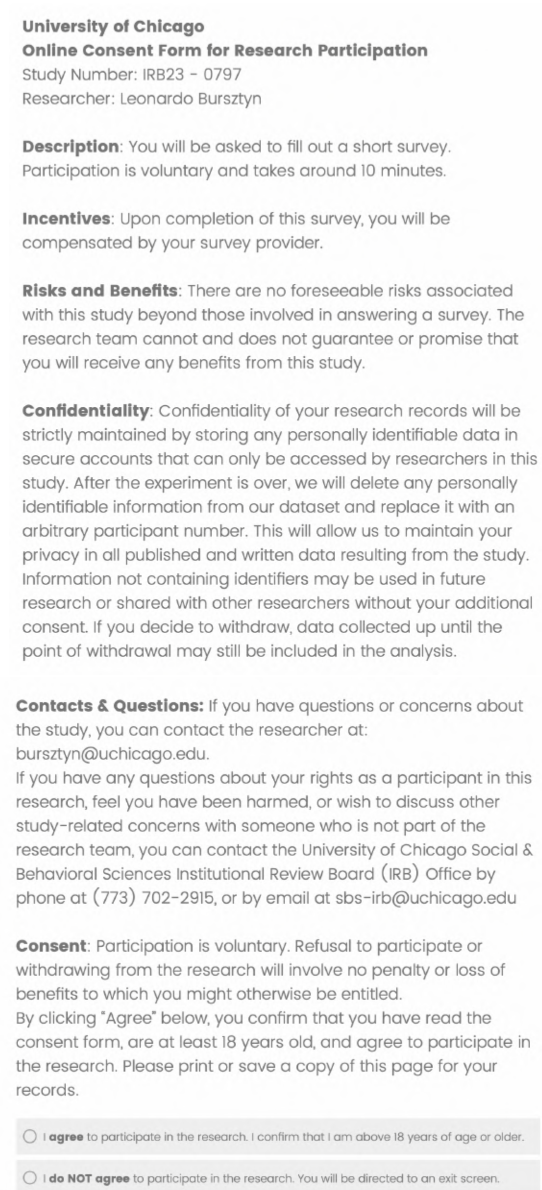

University of Chicago
Online Consent Form for Research Participation
Study Number: IRB23 - 0797
Researcher: Leonardo Bursztyn

Description: You will be asked to fill out a short survey. Participation is voluntary and takes around 10 minutes.

Incentives: Upon completion of this survey, you will be compensated by your survey provider.

Risks and Benefits: There are no foreseeable risks associated with this study beyond those involved in answering a survey. The research team cannot and does not guarantee or promise that you will receive any benefits from this study.

Confidentiality: Confidentiality of your research records will be strictly maintained by storing any personally identifiable data in secure accounts that can only be accessed by researchers in this study. After the experiment is over, we will delete any personally identifiable information from our dataset and replace it with an arbitrary participant number. This will allow us to maintain your privacy in all published and written data resulting from the study. Information not containing identifiers may be used in future research or shared with other researchers without your additional consent. If you decide to withdraw, data collected up until the point of withdrawal may still be included in the analysis.

Contacts & Questions: If you have questions or concerns about the study, you can contact the researcher at: bursztyn@uchicago.edu.
If you have any questions about your rights as a participant in this research, feel you have been harmed, or wish to discuss other study-related concerns with someone who is not part of the research team, you can contact the University of Chicago Social & Behavioral Sciences Institutional Review Board (IRB) Office by phone at (773) 702-2915, or by email at sbs-irb@uchicago.edu

Consent: Participation is voluntary. Refusal to participate or withdrawing from the research will involve no penalty or loss of benefits to which you might otherwise be entitled.
By clicking "Agree" below, you confirm that you have read the consent form, are at least 18 years old, and agree to participate in the research. Please print or save a copy of this page for your records.

- I **agree** to participate in the research. I confirm that I am above 18 years of age or older.
- I **do NOT agree** to participate in the research. You will be directed to an exit screen.

How frequently did you use each of the following social media platforms in the past month?

	Not at all	Once	Once a week	Twice a week	Every day
Instagram	○	○	○	○	○
Facebook	○	○	○	○	○
TikTok	○	○	○	○	○
Twitter	○	○	○	○	○

Please read the instructions of this survey carefully. We will ask you some questions to ensure your understanding of its content and give you a **bonus payment** if you answer correctly!

As mentioned, this survey is directed at university students' social media preferences. **We are studying how much students like you value Instagram.**

After this survey, we will conduct a study in which we will ask **students at your university to deactivate their Instagram accounts** for **four weeks** in exchange for a **monetary payment.**

"Deactivation" studies like this have been conducted in the past (e.g., by Mosquera et al., 2018 and Allcott et al., 2020) with close to 90% compliance.

Should students deactivate their Instagram accounts, they can go back to using it whenever they want, with their content and network unchanged, but they will then forgo any monetary payment.

To verify that students deactivate their Instagram accounts, we will **periodically visit their profiles and require them to upload screenshots of their app usage.**

The next part of the study involves asking you if you want to deactivate your Instagram account for four weeks in exchange for monetary payment.

For that we will need to collect your Instagram handle so that we can check whether you are active. We would also ask you to upload periodic screenshots of your phone's time use statistics if you are selected.

Would you be willing to

1. **deactivate your Instagram account for four weeks in exchange for monetary payment,**
2. provide your Instagram handle, and
3. provide screenshots of your phone's time use statistics?

○ Yes

○ No

How will we verify that selected users deactivate their Instagram accounts?

- ○ By periodically visiting their profiles and requiring them to upload screenshots
- ○ By asking them to install an app and requiring them to upload screenshots
- ○ By requiring them to upload screenshots
- ○ By periodically visiting their profiles

For how long will we ask selected users to deactivate their Instagram accounts?

- ○ Eight weeks
- ○ One week
- ○ Four weeks
- ○ Ten weeks

C.2.2 Step 0: Practice Good

Before we begin, we will ask you a series of hypothetical practice questions for you to get accustomed to our survey.

Uber

To understand how much people value their Uber account, **we ask university students, including you,** to decide whether to **deactivate their Uber account** for four weeks, in exchange for different monetary payments.

A computer will **randomly select one student from your university** to be eligible for the deactivation study.

Uber

Suppose the next questions involve real money. The computer will randomly generate an amount of money to offer you to participate in the deactivation study.

We will now ask you a series of questions offering you different payment scenarios in case you are selected for the deactivation study.

- If you accept any price scenario lower than the computer's offer, we will invite you to the deactivation study and give you the computer's offer.
- If you do not accept any price scenario lower than the computer's offer, we will not invite you to the deactivation study even if you are selected.

This rule means that the higher the amount you require the lower the chance that you will receive the computer's offer.

Therefore, while answering the following questions, **please choose carefully as each answer counts.**

We will now ask you a comprehension question based on the text above:

Which of the following statements is true?

- ○ The higher the amount I require the lower the chance that I receive the computer's offer
- ○ The higher the amount I require the higher the chance that I receive the computer's offer
- ○ The amount I require does not affect the chance that I receive the computer's offer

Uber

Over the next 4 weeks, would you like to take a break from ride-sharing apps?

Which of the following options would you prefer?

- ◯ **I take a break:** I deactivate my Uber account AND I receive $100
- ◯ **I do not take a break:** I keep my Uber account active

Uber

Over the next 4 weeks, would you like to take a break from ride-sharing apps?

Which of the following options would you prefer?

- ◯ **I take a break:** I deactivate my Uber account AND I receive $40
- ◯ **I do not take a break:** I keep my Uber account active

Uber

Over the next 4 weeks, would you like to take a break from ride-sharing apps?

Which of the following options would you prefer?

- ○ **I take a break:** I deactivate my Uber account AND I receive $80
- ○ **I do not take a break:** I keep my Uber account active

Uber

According to your answers to the previous questions, you would require a payment worth between $80 and $100 to deactivate your Uber account for four weeks.

Do you agree with the above statement about your **valuation**?

- ○ Yes
- ○ No

C.2.3 Step 1: Valuation Keeping Network

You have now completed the practice section. We will now ask you a series of questions about your **social media preferences**.

To establish appropriate payment amounts for the deactivation study, **we ask university students, including you**, to decide whether to deactivate their Instagram accounts for different monetary amounts.

A computer will **randomly select one student from your university** to be eligible for the deactivation study.

The next questions involve real money. The computer will randomly generate an amount of money to offer you to participate in the deactivation study.

We will now ask you a series of questions offering you different payment scenarios in case you are selected for the deactivation study.

Therefore, while answering the following questions, **please choose carefully as each answer counts.**

Instagram

Over the next 4 weeks, would you like to take a break from social media?

Which of the following options would you prefer?

○ **I take a break:** I deactivate my Instagram account AND I receive $120

○ **I do not take a break:** I keep my Instagram account active

Instagram

Over the next 4 weeks, would you like to take a break from social media?

Which of the following options would you prefer?

○ **I take a break:** I deactivate my Instagram account AND I receive $60

○ **I do not take a break:** I keep my Instagram account active

Instagram

Over the next 4 weeks, would you like to take a break from social media?

Which of the following options would you prefer?

◯ **I take a break:** I deactivate my Instagram account AND I receive $20

◯ **I do not take a break:** I keep my Instagram account active

Instagram

Over the next 4 weeks, would you like to take a break from social media?

Which of the following options would you prefer?

◯ **I take a break:** I deactivate my Instagram account AND I receive $40

◯ **I do not take a break:** I keep my Instagram account active

According to your answers to the previous questions, you would require a payment worth between $40 and $60 to deactivate your Instagram account for four weeks.

Do you agree with the above statement about your **valuation**?

- Yes
- No

C.2.4 Step 2: Valuation Removing Network

College Pulse has a panel exceeding 650,000 university students.* We are targeting universities with a high penetration of College Pulse.

We will now ask you to consider two additional options for a large-scale deactivation of Instagram at your university. **One of them will be randomly implemented** if we manage to recruit **more than two-thirds of the students at your university.**

We expect 90% of students to comply with deactivation based on previous studies (e.g., by Mosquera et al., 2018 and Allcott et al., 2020).

*See https://collegepulse.com.

Option 1 for the large-scale deactivation study:

In collaboration with College Pulse, we will ask students at your university sequentially whether they would like to deactivate their Instagram accounts.

Consider the scenario where we have asked **all participating students at your university** to deactivate their Instagram accounts for four weeks in exchange for a payment.

The next questions involve real money. The computer will randomly generate an amount of money to offer you to participate in the deactivation study.

We will now ask you a series of questions offering you different payment scenarios in case you are selected for the deactivation study.

Therefore, while answering the following questions, **please choose carefully as each answer counts.**

Over the next 4 weeks, would you like to take a break from social media if we ask all participating students at your university to take a break from social media in exchange for monetary payment?

If you had the choice to **also** deactivate your Instagram account, which of the following options would you prefer?

- ◯ **I take a break:** I deactivate my Instagram account if the other students deactivate their Instagram accounts AND I receive $60
- ◯ **I do not take a break**: I keep my Instagram account active if the other students deactivate their Instagram accounts

Over the next 4 weeks, would you like to take a break from social media if we ask all participating students at your university to take a break from social media in exchange for monetary payment?

If you had the choice to **also** deactivate your Instagram account, which of the following options would you prefer?

- ◯ **I take a break:** I deactivate my Instagram account if the other students deactivate their Instagram accounts AND I receive $20
- ◯ **I do not take a break**: I keep my Instagram account active if the other students deactivate their Instagram accounts

Instagram

Over the next 4 weeks, would you like to take a break from social media if we ask all participating students at your university to take a break from social media in exchange for monetary payment?

If you had the choice to **also** deactivate your Instagram account without payment, which of the following options would you prefer?

- ○ **I take a break:** I deactivate my Instagram account if the other students deactivate their Instagram accounts
- ○ **I do not take a break**: I keep my Instagram account active if the other students deactivate their Instagram accounts

Instagram

According to your answers to the previous questions, you would deactivate your Instagram account for four weeks without monetary payment, if we ask all participating students at your university to deactivate their Instagram accounts.

Do you agree with the above statement about your **valuation**?

- ○ Yes
- ○ No

C.2.5 Step 3: Product Market Valuation

Instagram

Option 2 for the large-scale deactivation study:

We know how much we would need to pay **every participating student at your university** to deactivate their Instagram accounts for four weeks.

To give everyone an equal chance to decide, we will randomly choose one of the students.

This student's identity will remain anonymous, and they can choose one from the following options:

1. **We ask all participating students with an Instagram account to deactivate it**, or
2. **We keep things as they are**.

If you decide for all participating students to deactivate their Instagram accounts:

- **We will pay the other students the amount they required** and we will **establish your payment, if any, below.**
- The deactivation study will be stopped for everyone only if you go back to using Instagram before the end of the four weeks.
- If the study is stopped early, you will not receive payment and we will pay the other students based on the actual time they spent in the study.
- If someone from the other participating students goes back to using Instagram before the end of the study, they themself will not receive any payment.

Instagram

Suppose that you decide for us to ask all participating students to deactivate their Instagram accounts:

Which of the following statements is correct?

- ◯ We will pay the other students what they required to deactivate their Instagram accounts
- ◯ We will pay the other students more than what they required to deactivate their Instagram accounts
- ◯ We will force the other students to deactivate their Instagram accounts

Instagram

The next questions involve real money. The computer will randomly generate an amount of money to offer you to participate in the deactivation study.

We will now ask you a series of questions with different payment scenarios in case you are selected.

Therefore, while answering the following questions, **please choose carefully as each answer counts.**

Instagram

Over the next 4 weeks, would you like students at your university to take a break from social media?

If you were not receiving payment, which of the following options would you prefer?

- ◯ **Everyone takes a break:** All participating students at my university deactivate their Instagram accounts
- ◯ **No one takes a break:** All participating students at my university keep their Instagram accounts active

Instagram

Over the next 4 weeks, would you like students at your university to take a break from social media?

Which of the following options would you prefer?

- ◯ **Everyone takes a break:** All participating students at my university deactivate their Instagram accounts
- ◯ **No one takes a break:** All participating students at my university keep their Instagram accounts active AND I receive $200

According to your answers to the previous questions, you would forgo a payment above $200 to have all participating students at your university, including you, deactivate their Instagram accounts for four weeks.

Do you agree with the above statement about your **valuation**?

- ◯ Yes
- ◯ No

C.2.6 Qualitative Questions

What is the percent chance that we will recruit more than two-thirds of the students at your university?

Enter a number between 0 and 100.

Would you prefer to live in a world with or without Instagram?

- ◯ I would prefer to live in a world with Instagram
- ◯ I would prefer to live in a world without Instagram

You mentioned you would prefer to live in a world without Instagram.

Why do you still use it?

Please rank the following options from your most preferred (1) to your least preferred (3).

- I deactivate Instagram and every other student at my university keeps using it
- Every student at my university, including me, deactivates Instagram
- No one deactivates Instagram

How would you feel if you were the only one who deactivated Instagram and everyone else kept using it?

What fraction of your mutual friends on Instagram are fellow college students? Please enter your response in percent.

Enter a number between 0 and 100.

To what extent do you agree with the following statements?

	Strongly disagree	Somewhat disagree	Neither agree nor disagree	Somewhat agree	Strongly agree
I would worry about missing out on new trends or memes on social media if I were to quit it.	○	○	○	○	○
I use social media immediately upon waking up in the morning.	○	○	○	○	○
I feel pressure to share or post interesting content on social media.	○	○	○	○	○
I feel anxious if I haven't checked social media for a while.	○	○	○	○	○
I would worry that I would be socially isolated if I quit social media.	○	○	○	○	○
I feel left out when I see posts or videos about gatherings or events on social media that I wasn't part of.	○	○	○	○	○

How frequently did you share content from social media with your friends in the past month?

	Not at all	Once	Once a week	Twice a week	Every day
Instagram	○	○	○	○	○
Facebook	○	○	○	○	○
TikTok	○	○	○	○	○
Twitter	○	○	○	○	○

Now some **demographic** questions.

Which of the following describes you more accurately?

- ◯ Male
- ◯ Female
- ◯ Other / Prefer not to say

What is your age?

Thank you for completing the survey! Your incentive will be delivered to you within 2-3 business days. Remember to check your spam folder. If you don't receive it, contact support@collegepulse.com. Your feedback is highly appreciated!

C.3 Luxury and Vintage Good Survey

University of Chicago
Online Consent Form for Research Participation
Study Number: IRB23 - 1227
Researcher: Leonardo Bursztyn

Description: You will be asked to fill out a short survey. Participation is voluntary and takes around 5 minutes.

Incentives: Upon completion of this survey, you will be compensated by your survey provider.

Risks and Benefits: There are no foreseeable risks associated with this study beyond those involved in answering a survey. The research team cannot and does not guarantee or promise that you will receive any benefits from this study.

Confidentiality: Confidentiality of your research records will be strictly maintained by storing any personally identifiable data in secure accounts that can only be accessed by researchers in this study. After the experiment is over, we will delete any personally identifiable information from our dataset and replace it with an arbitrary participant number. This will allow us to maintain your privacy in all published and written data resulting from the study. Information not containing identifiers may be used in future research or shared with other researchers without your additional consent. If you decide to withdraw, data collected up until the point of withdrawal may still be included in the analysis.

Contacts & Questions: If you have questions or concerns about the study, you can contact the researcher at: bursztyn@uchicago.edu. If you have any questions about your rights as a participant in this research, feel you have been harmed, or wish to discuss other study-related concerns with someone who is not part of the research team, you can contact the University of Chicago Social & Behavioral Sciences Institutional Review Board (IRB) Office by phone at (773) 702-2915, or by email at sbs-irb@uchicago.edu

Consent: Participation is voluntary. Refusal to participate or withdrawing from the research will involve no penalty or loss of benefits to which you might otherwise be entitled.
By clicking "Agree" below, you confirm that you have read the consent form and agree to participate in the research. Please print or save a copy of this page for your records.

- ○ I **agree** to participate in the research.
- ○ I **do NOT agree** to participate in the research. You will be directed to an exit screen.

Do you currently own products from luxury brands that you purchased **yourself**? Please tick all that apply.

- ☐ Balenciaga
- ☐ Rolex
- ☐ Yves Saint Laurent (YSL)
- ☐ Gucci
- ☐ Burberry
- ☐ Tiffany & Co.
- ☐ Versace
- ☐ Givenchy
- ☐ Louis Vuitton
- ☐ I do not own products from any luxury brands
- ☐ Chanel
- ☐ Other:
- ☐ Swarovski

Do you prefer to live in a world with or without any luxury fashion brands (such as Louis Vuitton, Gucci, Chanel, Yves Saint Laurent (YSL), Balenciaga, Versace, Rolex, Tiffany & Co., Burberry, Givenchy, and Swarovski)?

- ○ I prefer to live in a world with luxury fashion brands
- ○ I prefer to live in a world without any luxury fashion brands

Do you currently own an IPhone that you purchased yourself?

- ○ Yes
- ○ No

Which IPhone model do you own?

Do you prefer to live in a world where Apple releases the IPhone every year or every other year?

- ◯ I prefer to live in a world where Apple releases the iPhone **every other year**
- ◯ I prefer to live in a world where Apple releases the iPhone **every year**

Are you planning to buy or have you already bought the new IPhone 15 this year?

- ◯ Yes
- ◯ No

Do you prefer to live in a world with or without any IPhones?

- ◯ I prefer to live in a world with IPhones
- ◯ I prefer to live in a world without any IPhones

Which of the following describes you more accurately?

- ◯ Male
- ◯ Female
- ◯ Other / Prefer not to say

What is your age?

What was your TOTAL household income, before taxes, last year?

- ◯ $0 - $20 000
- ◯ $20 000 - $50 000
- ◯ $50 000 - $90 000
- ◯ $90 000 - $150 000
- ◯ $150 000 - $200 000
- ◯ $200 000+

Which category best describes your highest level of education?

- ◯ Some High School
- ◯ High School Degree
- ◯ Some College
- ◯ College Degree
- ◯ Master's Degree
- ◯ Doctoral Degree

What racial or ethnic group best describes you?

- ◯ White
- ◯ Black or African-American
- ◯ Hispanic or Latino
- ◯ Asian or Asian-American
- ◯ Middle Eastern
- ◯ Other